

Delphix Engine 6.0.13.0 Documentation

Documentation 6.0.13.0

Exported on 08/16/2023

Table of Contents

Release Notes	6
Release Notes 6.0.x	7
Data Source Integration (Plugin) Release Notes	208
Overview of Data Virtualization	389
Delphix Engine Overview	390
Getting Started with Data Sources	391
Delphix Product Information	392
Glossary of Major Delphix Concepts	395
Product Icon Reference	404
Deployment	406
Standard Deployment Architecture	407
Checklist of Information Required for Installation and Configuration	408
Network Connectivity Requirements	413
Installation and Initial System Configurations	418
Validating Host Deployment with HostChecker	428
Deployment for VMware	429
Deployment for KVM	436
Deployment for Hyper-V	440
Deployment for AWS EC2	446
Deployment for Microsoft Azure	456
Deployment for Google Cloud Platform	465
Deployment for OCI	471
Deployment for IBM Cloud	480
Hotfix Information	488
Configuration	489
Registration Management	490
User and Authentication Management	492
Network and DNS Management	537
NFSv4 Configuration	572
Capacity and Resource Management	574
Monitoring and Log Management	606
Performance Analytics Management	683

Usage Data Management.....	733
Starting, Stopping, and Restarting Your Engine.....	735
Introduction to Privilege Elevation Profiles.....	738
Upgrade	740
Upgrading the Delphix Engine: Overview	741
Upgrade Prerequisites	743
Downloading the Upgrade Image	745
Uploading the Upgrade Image	746
Applying the Upgrade	757
Post Upgrade.....	760
Security.....	761
Security Principles	762
Product Security	763
Object Security.....	813
System Configuration	818
GUI Security.....	819
Masking Sensitive Data.....	820
Audit Logs.....	821
Support Security	822
Password Policies	824
Additional Topics	826
Datasets.....	828
Getting Started.....	829
Oracle Environments and Data Sources	903
SQL Server Environments and Data Sources.....	1152
SAP ASE Environments and Data Sources	1325
IBM Db2 Environments and Data Sources	1436
PostgreSQL Environments and Data Sources	1538
HANA Environments and Data Sources	1618
Unstructured Files and App Data	1684
Virtualizing Oracle E-Business Suite	1726
Best Practices.....	1814
Introduction to Delphix Architecture	1814
Hypervisor and Host	1814

Network	1814
Storage	1814
Data Protection	1814
Source DB and OS Settings.....	1814
Target DB and OS Settings.....	1814
Staging Target.....	1815
Architecture Checklist (Consolidated Best Practices).....	1815
Architecture Checklist FAQ	1815
Architecture Checklist FAQ	1816
Architecture Checklist.....	1821
Best Practices for Hypervisor Host and VM Guest	1832
Best Practices for Network Configuration	1836
Best Practices for Storage	1839
Best Practices for Delphix Engine Data Protection	1841
Best Practices for Source DB and OS Settings.....	1844
Best Practices for Target DB and OS Settings.....	1845
Best Practices for Staging Targets	1850
Data Backup and Recovery Solutions	1852
Delphix Data Vault.....	1853
Backup and Recovery Strategies for the Delphix Engine	1859
Delphix Replication	1869
Selective Data Distribution	1911
Delphix Self-Service	1925
Delphix Self-Service Admin Guide.....	1926
Delphix Self-Service Data User Guide	1986
Delphix Masking	2028
Developer's Guide	2029
Command Line Interface Guide	2030
Web Services API Guide	2304

Here you will find all the information you will need to use the Delphix Continuous Data Engine for data virtualization. Learn how to deploy our application, use our features, or tune it for optimal performance. We have organized this content into several categories that you can browse via the navigation bar.

Release Notes

 Unknown macro: 'navbox'

Release Notes 6.0.x.

Welcome to the 6.0.x release of the Delphix Engine database virtualization system.

- [New Features](#)
- [API Changes](#)
- [Support Matrices](#)
- [Migration and Compatibility](#)
- [Fixed Issues \(6.0.x\)](#)
- [Known Issues](#)
- [Deprecated and End-of-Life Features](#)
- [Licenses and Notices](#)

New Features

This section describes the new features and enhancements on the Delphix products.

Release 6.0.13.0

What's New in this Release

- **Staging Push for SQL Server**
This release supports staging push for SQL Server. Staging push allows you to push data to Delphix if you have tools, configurations, or security concerns that do not work with the existing pull model. Staging push enables, for example, compatibility with data stored in backup systems and gives you control over the staging database. For more information, see [Staging Push Implementation for SQL Server](#).
- **Support for TLS 1.3**
This release provides support for TLS 1.3 for connections from the virtualization engine to host and engine-to-engine communication.
- **Automatic vPDB restart for Oracle Single Instance vPDBs in a Linked CDB**
This release now supports the automatic vPDB restart feature for Oracle Single Instance vPDBs (12.2 and higher) that are provisioned into a Single Instance Linked CDB. Auto-restart allows the engine to detect and restart Single Instance vPDBs if the remote host has been restarted. Users can choose to enable or disable automatic restart on a per-vPDB basis.
- **Transparent Data Encryption (TDE) for Oracle Multitenant RAC (Oracle 18c and 19c)**
The Oracle TDE feature encrypts the sensitive data (database tables and tablespaces) stored on the disk. This prevents misuse of the data if the disks or storage mediums are lost or stolen. The data is transparently decrypted for authorized users when they access the data. Our large enterprise customers leverage Oracle RAC configurations for their business-critical applications. In the 6.0.12.0 release, Delphix introduced support for provisioning a TDE-encrypted vPDB to a linked (Physical CDB) target RAC environment on Oracle 18c and 19c only. This release adds support for Oracle 12.2 and 21c as well. For more information, see [Provisioning a TDE-enabled vPDB to a Cluster Target](#).
- **Db2 Plugin Version 4.1.3**
This plugin release resolves some issues that are listed in the [Db2 Release Notes](#).
- **EBS Plugin Version 3.0.3**
This plugin release resolves some issues that are listed in the [EBS Release Notes](#).

Certifications

- **Virtualization**
- Microsoft Windows Server 2022
- Oracle E-Business Suite (EBS) 12.2.11 in OCI



- For DSIL plugin version-specific release notes, see [Data Source Integration \(Plugin\) Release Notes](#).
- For masking release notes, see [Masking Release Notes](#).
- For DCT release notes, see [DCT SaaS](#) or [DCT Multicloud](#).

Release 6.0.12.0

What's New in this Release

- **HANA Staging Push**
The HANA 2.0 plugin introduces the Staging Push feature. This feature provides a new data ingestion

mechanism that helps users to push data into the Delphix provided mount point on their own. For more information, see [Delphix Architecture for HANA](#).

- **Multiple Device Removal**

You can now remove multiple storage devices at a time when engines are over-provisioned or when moving to new storage. This is done only after ensuring sufficient resources to support the removal.

- **Transparent Data Encryption (TDE) for Oracle Multitenant RAC (Oracle 12.2 and 21c)**

The Oracle TDE feature encrypts the sensitive data (database tables and tablespaces) stored on the disk. This prevents misuse of the data if the disks or storage mediums are lost or stolen. The data is transparently decrypted for authorized users when they access the data. Our large enterprise customers leverage Oracle RAC configurations for their business-critical applications. We have added support for provisioning a TDE-encrypted vPDB to a linked (Physical CDB) target RAC environment on Oracle 18c and 19c in this release. For more information, see [Provisioning a TDE-enabled vPDB to a Cluster Target](#).

Certifications

- Oracle 21c on SLES 15 SP3 (6.0.11.0+)
- Oracle 19c on SLES 15 SP3 (6.0.11.0+)
- Oracle 21c on SLES 15 SP2 (6.0.11.0+)
- Oracle 19c on SLES 15 SP2 (6.0.11.0+)
- AppData on RHEL 8.4 (6.0.11.0+)



- For masking release notes, see [Masking Release Notes](#).
- For DCT release notes, see [DCT SaaS](#) or [DCT Multicloud](#).

Release 6.0.11.0

What's New in this Release

- **OAuth2 API Support**

The Virtualization and Masking engine APIs are now accessible via OAuth2 tokens that improve Delphix's security offerings. For more information, see [Configuring OAuth2 Authentication for API Access](#).

- **Oracle Support for Exadata, Exadata Cloud, or Exadata Cloud-at-Customer (ExaCC) Cluster**

This release adds support for Exadata, Exadata Cloud, or Exadata Cloud-at-Customer (ExaCC) Cluster for Oracle databases. For more information, see [Oracle Support Matrix](#).

- **Apply Hotfixes using the Self-Service UI**

You can now apply hotfixes using the self-service upgrade UIs.

- **TDE for Oracle Multitenant - Support for Rekey**

Oracle Advanced Security Transparent Data Encryption (TDE) provides the ability to create virtual pluggable databases with a new key (independent from the source database). This is to facilitate another layer of security - ensuring that different keys are used in the production and non-production systems. For more information, see [Provisioning a TDE-enabled vPDB](#).

- **Modify DBID for Non-multitenant Oracle VDBs after Provisioning**

You can now generate a new DBID after a VDB is provisioned and refreshed. For more information, see [Generate a New DBID for Oracle VDBs](#). Support for multitenant databases is currently not available.

- **New Wizard for Creating Replication Profile**

A new multi-step wizard is now available that replaces the previous in-place editing option to manage Replication Profiles. For more information, see [Replication User Interface](#).

Certifications

- Oracle 21c on RHEL 8.4, RHEL 8.3, SLES 15 SP1 (6.0.11.0+)

- Oracle 19c on RHEL 8.4 (6.0.10.0+)
- Exadata Cloud-at-Customer(ExaCC)/Exadata support for Oracle 12.2 on OEL 7.8 and OEL 7.9



- For masking release notes, see [Masking New Features](#).
- For DCT release notes, see [DCT SaaS](#) or [DCT Multicloud](#).

Release 6.0.10.0

Virtualization

- **SQL Server-Support for Azure Storage Backups**
SQL native backups can now be read directly from Azure Cloud storage. For more information, see [Restoring SQL Backups Stored in Azure Cloud Storage](#).
- **Support dSource Upgrades From Non-MT to MT in Oracle**
Once a dSource is converted to a multitenant PDB, you will be able to share its storage blocks with its non-multitenant predecessor. Delphix will only store the incremental changes to the database. For more information, see [Prepare and Upgrade a Non-MT Oracle dSource to MT](#).
- **Oracle-Provision to the Latest Point-In-Time on RAC**
Provision to the latest point-in-time is now supported.
- **Flexible ORACLE_HOME Permissions Configuration**
Removed the need to set permissions of the `$ORACLE_HOME/dbs` subdirectory using STARTUP SPFILE syntax to simplify Oracle operations. For more information, see [Requirements for Oracle Hosts and Databases](#).
- **Support for Manually Starting an Oracle VDB**
An Oracle VDB can now be manually started. For more information, see [Manually Starting a VDB](#).
- **Attach and Detach Oracle CDB Containers**
Detaching, attaching, and linking of the Oracle CDB containers is now supported via CLI. For more information, see [CLI Cookbook: Attaching, Detaching, or Linking a CDB](#).
- **Password Vaults and Remote Hooks for UI**
In the 6.0.9.0 release, we introduced the ability to use password vaults with hooks. This allows our customers to ensure a high level of security with all operations with external systems. This can now be configured via the user interface. For more information, see [Passing Credentials Securely to Hook Operations](#).
- **NFSv4**
NFSv4 is now enabled by default. For more information, see [NFSv4 Configuration](#).
- **12-month Support for Upgrades and Forward Compatible Replication (FCR)**
To better align with the Delphix support program, Engine upgrades and FCR operations must be within versions no more than 12 months apart. For example, upgrading to version 6.0.10.0 will require the previous version to be at least 6.0.4.0. For more information, see [Upgrade Matrix](#) and [Replication Overview](#).
- **HANA Plugin Port Control**
You can now keep the port numbers consistent throughout the HANA VDB life cycle so that the connections made to the VDBs are not disrupted during their life cycle. For more information, see [Provisioning HANA VDBs: An Overview](#).
- **User-Specified Mount Path For Db2 dSource**
You can now specify a mount path of your choice to host the dSource dataset on the target host. For more information, see [Linking a Db2 dSource](#).
- **Staging Push Automation for Db2 dSource**
You can now use a set of scripts that can be used to automate the restore and rollforward operations on the dSource. For more information, see [Staging Push Automation](#).

Certifications

Virtualization

- [OCI VM.Standard.E4](#)
- [AWS r5n.24xlarge](#)
- [Oracle 21C RH 8.3](#)

Release 6.0.9.0

Virtualization

- **Delphix Data Vault - Additions**
In continuation to the **Delphix Data Vault for ransomware protection** released in 6.0.8.0, this release enables you to manage this feature from the Delphix Engine user interface and also lets you monitor the regular valid replication receives. For more information, see [Delphix Data Vault](#).
- **TLS 1.3**
The TLS 1.3 support is added as an available secure connection option at the Engine Admin Console to be used between engines.
- **Phonehome Data Collection Frequency**
The default collection period for the Phonehome users is now changed to daily. For more information, see [System Configurations - Enable Phone Home](#).
- **Password Vaults for Remote Hooks**
With this release, the hooks running on environments can now obtain credentials from the engine and its configured password vaults. These credentials can be used to perform custom authentication tasks in a secure manner. Currently, this feature is supported via the command line interface only. For more information, see [Passing Credentials Securely to Hook Operations](#).
- **Delphix Integrations (dxi) Docker Image**
A new dockerized version of the dxi library is now available. For more information, see [Docker Image](#).

Certifications

- **Virtualization**
 - [EBS 12.2 on Oracle 19c](#)
 - [Db2 11.1.4.6 Fix pack](#)
- **Hypervisors/Clouds**
 - [ESXi 7.0 U2](#)

Release 6.0.8.0

Virtualization

- **Delphix Data Vault**
The Delphix Data Vault for ransomware protection (accessible via CLI) enables organizations to recover access to their application data much faster than traditional backup solutions after malicious attacks. It relies on the new **Data Vault** Replication feature, which replicates critical business DB data stored on Delphix engines to a new target engine called Data Vault. Once securely stored on the Data Vault, the replicated DB data can be used to recover business applications upon a ransomware attack with very low RTO and RPO.
- **Dxi Executable and Support for Encrypted Credentials**
We will be distributing the dxi CLI, a Delphix solution built to facilitate simpler and seamless integration of Delphix Platform Operations into existing workflows, as Windows, macOS and RHEL binaries. This will

simplify adoption and remove the requirement on Python. Delphix has also adding encryption for the login credentials.

- **Expansion of Retention Period on Replicated Objects**

At present, when snapshots on the replication source engine are deleted (either due to retention policies or user action), the next replication job will delete those snapshots on the replication target engine. This improvement will allow you to extend the retention period of replicated objects on target engines while keeping the original retention at the source. Once the object in the target engine reaches its retention period, it will be flagged and deleted by the policy agent based on a daily schedule.

- **Db2 Staging Push**

Delphix now supports Staging Push for non-DPF Db2 databases. The Staging Push architecture will allow organizations to bring their data to Delphix, with their own tools and standard processes. This facilitates the use of any backup tool, a major ongoing source of requests. This should dramatically increase the volume of data that can be easily managed by Delphix.

- **HANA Plugin Staged Architecture**

Delphix introduces a staging architecture for HANA virtualization. This will make it consistent with other virtualized data sources. This new architecture will build a foundation for future staging push capabilities, as Delphix has begun to introduce for other platforms. These changes, together, will allow us to support more prospective organizations with various SAP-certified, 3rd-party backup applications for HANA. Delphix will continue to support the pull ingestion method with HANA native backups and logs.

- **ASE Native Encryption Support**

SAP ASE version 16.0 introduces the ability to fully encrypt databases, provides protection for all the data, indexes, and transaction logs in a database. This offers full database protection while allowing the user to query and manage the data as usual, as the encryption is transparent to existing functions. In response to customer demand, we have added a security enhancement to support encrypted ASE databases. This allows customers to maintain the ASE encryption that is active on their sources and propagate that through to their VDBs.

- **System Tunable Interface**

A new web service API is introduced which allows you to set or receive the values of a system tunable via Delphix CLI. A Support engineer can now provide context via Support case to modify these values.

- **Oracle Customized Full Backups**

There is a rarely-seen bug in Oracle that results in some blocks not being written to the datafiles during an Oracle SnapSync operation. When this happens, the datafiles can become incomplete and provisioning/refreshing from that snapshot might fail. We are providing a SnapSync option via CLI that you can customize to accept all datafiles during an Oracle SnapSync operation to prevent this error.

Release 6.0.7.0

Virtualization

- **Simplified Connection Management for Oracle databases**

This feature streamlines the way that Delphix communicates and interacts with Oracle databases by simplifying the connection management infrastructure. Prior to this release, connections were established to Oracle databases using two different methods (remote connections from the Delphix Engine and local connections from the Delphix toolkit) and communication was performed with Oracle databases using two different users (a Delphix OS user and a Delphix DB user). Starting with this release, all communication with Oracle databases will be performed locally on the Oracle host and all connections to Oracle databases will be established using OS authentication. Existing Oracle dSources and VDBs will continue to function with no user intervention required. This feature results in several key benefits for Oracle DB customers such as elimination of the requirement for a Delphix DB user when linking, automated PDB discovery, elimination of Delphix interaction with any network listener, and many more.

- **Virtualization SDK Support for Password Vaults**

Building off of the existing CyberArk and Hashicorp support for Oracle, SQL Server, and SAP ASE database user credentials, Delphix Virtualization will extend password vault coverage to the virtualization SDK (vSDK).

This will enable data sources that are connected via a vSDK plugin to incorporate this more secure method of authentication.

- **SAP ASE Device Mapping Improvements**

The Delphix Virtualization experience with SAP ASE heavily relies on and mirrors a database’s device allocation from the initial load (creating the dSource) to provision (creating VDBs). As these source device allocations shift over time, Delphix maps these changes and propagates them to their associated Delphix objects. However, dramatic device layout changes can negatively impact performance. This enhancement provides a quality-of-life (QoL) improvement to the overall SAP ASE experience by providing better error handling and escape valves should a dSource get into a bad state due to a major device layout shift.

- **Improved Storage Utilization for Large Pools**

Up through the 6.0.6 release, Delphix Virtualization has enforced a storage usage limit of 85%. Once met, this limit will cause certain API operations to be disabled to ensure engine data integrity. In the 6.0.7 release, this threshold is relaxed significantly. The new thresholds are as follows:

- “Warning”: when 85% of the total storage quota is reached or 1536GB of free storage is remaining (whichever is less), which can be resolved/ignored, with no impact on system behavior.
- “Critical”: when 90% of the total storage quota is reached or 1024GB of free storage is remaining (whichever is less), which cannot be resolved/ignored, with some impact on system behavior.
- “Minimum”: when 95% of the total storage quota is reached or 512GB of free storage is remaining (whichever is less). In this case, a critical fault is raised and cannot be resolved/ignored, with a substantial impact on system behavior (stop policies, VDB operations, etc).

- **PVSCSI support**

In addition to LSI Logic, with the 6.0.7 release, Delphix has added support for the VMware Paravirtual vSCSI controller (aka PVSCSI). While VMware designed PVSCSI to support very high throughput with minimal processing cost, the performance improvements on Delphix engines can vary from case to case. In 6.0.7, we also support manual changes from LSI Logic to PVSCSI for currently deployed engines.

Release 6.0.6.0

Virtualization

- **Solaris x86 to Linux x86 Oracle DB Provisioning**

This feature allows provisioning of Oracle Virtual Databases from Solaris x86 dSources to Linux x86 target environments.

- **TDE for Oracle Multitenant**

Oracle Advanced Security Transparent Data Encryption (TDE) provides the ability to encrypt sensitive application data on storage. Delphix will now support TDE for Oracle 12cR2, 18c and 19c multitenant. This release introduces support for single-instance linked container databases (CDBs) using software keystores. Virtual Container Database (vCDBs), RAC and rekeying of the TDE encryption keys are not supported in this release.

-  Please note the following important restrictions for **TDE for Oracle Multitenant** feature:

 - TDE-enabled vPDBs must be provisioned to a linked CDB, not a vCDB.
 - RAC dSources and target CDBs are not supported.
 - The Oracle version must be 12.2 or higher (12.1 is not supported).
 - System tables or tablespaces either in the PDB or CDB must not be encrypted.
 - Oracle Key Vault is not supported.
 - Hardware keystores are not supported.
 - Keystores must not be on ASM storage.
 - The dSource from which the initial provision is done must be encrypted when it is linked. Existing dSources cannot be encrypted without unlinking and creating a new dSource.

- Encrypting an already-provisioned unencrypted vPDB (with clear data) which is managed by Delphix is not supported.

- **Single to Multitenant VDBs**

Oracle announced the end of support for non-multitenant databases in their 20c release, and as such, Oracle DB customers are planning their upgrade and migration programs. Delphix will now support provisioning a virtual pluggable database from a non-multitenant virtual database.

- Added support for **HashiCorp namespaces**.

Certifications

- **Virtualization**

- ASE 15.7/16 on RHEL 7.9
- Oracle 12.1 on RHEL 7.9 and SLES 12 SP5
- Oracle 12.2 on RHEL 7.9 and SLES 12 SP5
- Oracle 19c on RHEL 7.9 and SLES 12 SP5

- **Hypervisors/Clouds**

- ESXi 7.0 U1

Release 6.0.5.0

Virtualization

- **NFSv4 Support:**

Support has been added for Oracle and ASE on AIX.

- Expanded Replication: Replication of non-data objects

Our customers are increasingly using replication to facilitate moves of data across network boundaries, to the cloud and for DR purposes. We've had long-standing requests to replicate more than just the data, and in this release, we will support replication of users, roles, permissions, policies and configuration templates.

- Upgraded Windows Connector:

The Windows Connector will now support newer versions of Microsoft's .NET framework (4.x), which encompasses myriad higher security standards, new functionality, etc. Previously, the connector relied on .NET 3.5 due to two dependencies: SQL Server and Powershell, both have since been removed with SQL Server 2016+ and Powershell update in 6.0.3.0.

- Db2 Extensible Ingestion:

We will now support an extensible model for ingesting Db2 data. In this new, additive, model, we will support customers manually performing a restore & roll forward of their staging database to Delphix from native backups or arbitrary third-party backup tools which integrate directly with Db2. This will allow customers to bring data from whatever system or backup they have and restore to an exact point in time, as needed.

Certifications

- **Virtualization**

- ASE 16.0 on RHEL8.1 and RHEL8.2 on 6.0.4+
- ASE 16.0 on SLES12.4 on 6.0.4+
- Oracle 19.7 on RH7.8 and RH8.0 on 5.3.9+ and 6.0.3+
- ESX 7.0
- NFS v4 support on AIX
- IBM Cloud Catalog. Delphix is now available in the IBM Cloud Catalog, a private marketplace for trusted IBM Technology partners that is offered to large IBM enterprise customers. In 6.0.5 we will

start with a few certified instances for virtualization and masking and will grow our presence as more the business justifies the cost and efforts. Specifically, we support the following instances:

- mx2-8x64
- mx2-16x128
- mx2-32x256
- mx2-48x384
- Oracle Cloud: The following are newly supported instance types:
 - VM.Standard2.8
 - VM.Standard2.16
 - VM.Standard2.24

Release 6.0.4.0

Virtualization

- **HashiCorp and expanded CyberArk Support:**

Delphix has extended both CyberArk and HashiCorp Vault support to Oracle Database Users in addition to previously supported ASE and MSSQL domain users. GUI support for HashiCorp Vault has been added during setup to authenticate host users and database users.

 The HashiCorp namespace Enterprise feature is supported starting 6.0.6.0.

- **NFSv4 Support:**

Support has been added for SuSE and Db2 on AIX.

Certifications

- **Virtualization**

- OCI Support
- NFSv4 support for Db2 on AIX and SuSE
- SQL Server Instances with a Managed Service Account

Release 6.0.3.0

Virtualization

- **CyberArk and Hashicorp Vault Support for Virtualization:**

Delphix is introducing password vault support to authenticate environment and database linking and will support both CyberArk with Oracle and Hashicorp with CLI only.

- **Capacity Management:**

Understanding where and how storage is used on Virtualization Engines can be a challenge, in particular, understanding how and where space is held and how to recover it. In this release, Delphix provides better details of held space, particularly around locked objects, and provides clear instructions about what steps are required to free up space.

- **Diagnosability:**

Additional performance health-check analytics in phone-home have been added to better troubleshoot and understand customer problems.

- **Powershell Upgrade:**

Delphix is reducing our requirements for Windows hosts running PowerShell by allowing you to use any PowerShell version from 2.0 to 5.1. Delphix will now use the default available PowerShell version on each host. When specifying hooks (such as “configureClone”), users may specify whether to use 2.0 or whatever PowerShell version is installed on the host.

- **Support for Oracle Read-only Homes:**

Delphix is introducing support for Oracle read-only homes, which is a new Oracle feature starting with Oracle 18c. In a read-only Oracle home, all the configuration data and log files reside outside of the read-only Oracle home. This feature allows you to use the read-only Oracle home as a software image that can be distributed across multiple servers.

- **Replication Performance:**

Delphix will continue to improve replication performance for replication specifications that include multiple objects and single-object replication throughput.

- **SAP ASE Support for VDB Upgrade:**

SAP ASE Customers will now be able to validate DBMS Upgrades with this feature that enables provisioning VDBs to a higher version than the source DB (e.g. ASE 15.7 > ASE 16).

- **Shared NFS for Toolkits**

With this release, Delphix introduces shared NFS for clustered environments. Customers wish to use a common NFS mount point, in which the Delphix toolkit for each cluster node can be deployed. The product today only creates a directory with appliance UUID and OS user in the folder name and uses this for detection to determine if a host is already managed by that Engine. As such, this prevents the customer from utilizing common NFS storage due to name conflict.

When a new environment is created, upgraded or if the toolkit path is changed, a new toolkit is created with naming convention Delphix_COMMON_<last 12 characters of engine ID>_<last 12 characters of host ID>_<user ref>_<host or cluster> for common directory and Delphix_<last 12 characters of engine ID>_<last 12 characters of host ID>_<user ref>_<host or cluster> for user directory.

With this change, the customer can use mounts on shared file systems (like NFS) as a toolkit path for clustered environments without any naming conflict.

 This change is not intended for windows environments.

Certifications

- **Virtualization**

- EBS 12.2 with RHEL 7.6
- PostgreSQL 12.1 & 12.2 with RHEL 7.8
- Oracle 11g R2 and Oracle 19c with RHEL 7.8 on 5.3.9.0 and 6.0.2+

Release 6.0.2.0

Virtualization

- **Support for Db2 Database Partition Feature (DPF):**

Delphix has long supported distributed Db2 (running on Unix/Linux Systems). However, Db2 supports partitioned databases as a means of scaling to larger, more complex systems. With this release, Delphix will now support Db2 DPF allowing you to scale to an increasing number of your Db2 databases.

- **Windows Authentication for SQL Server:**

You will now be able to use Windows Authentication to link SQL Server databases. Rather than providing both a database user and a Windows user to ingest data, you can leverage one set of credentials (a Windows OS user) to perform all source operations. This capability will simplify SQL Server deployments and reduce Delphix's security requirements on source databases.

- **Smart Failover:**

Smart Failover allows the Delphix Administrator to simplify failover processes by automating object conflict resolution. By selecting a new option “Automate Object Conflict Resolution” before the failover process starts, the failover process will rename all conflicting objects and show a report of all object changes at the end.

- **NFSv4 Support:** In 6.0.2 Delphix will start providing NFSv4 for data sources running on RedHat 7.0 or later. NFSv4 can be enabled using the CLI. Support for additional host OS versions will be added in subsequent releases. Delphix will consider enabling NFSv4 by default for those supported configurations in a future release.
- **Support Bundles not Required for Upgrade:** When upgrading from 6.0.0 or greater to a release 6.0.2 or greater, we no longer require support bundles to be sent to Delphix. This allows you to execute more self-service upgrades.

Certifications

- ASE 16.0 on AIX 7.1
- AWS r5n Instance Support: r5n.2xlarge, r5n.8xlarge, r5n.16xlarge
- Azure E Series Instance Support: E8s_v3, E16s_v3, E32s_v3
- Masking support of Oracle 19c.

Release 6.0.1.0

- **Masking Extended Connectors:** A very common request for masking has been to support additional data sources, outside of [the currently supported list](#). Thus, the next step in the strategy is the release of Masking Extended Connectors, which will allow our customers to add JDBC drivers to the masking engine to facilitate the masking of additional data sources. This will allow masking to be used for other common databases that can be accessed via JDBC, like SAP HANA, Informix, etc.
- **SQL Server CDC Support:** We have expanded our support for SQL Server databases using [Change Data Capture](#) (CDC), a SQL Server feature that captures all the change information that is applied to the databases and stores it in change tables. Now, users will have the ability to preserve CDC data and enable CDC for SQL Server VDBs.

Certifications

- **Virtualization**
 - ASE 16 and 15.7 with Solaris SPARC 11U3 and SPARC 11U4
 - ASE 16 and 15.7 with RHEL 7.7
 - SQL Server 2019 Support with Windows 2016 and Windows 2019
 - Oracle 19c with SUSE SLES 15 SP1
 - Oracle 19c with Solaris 11 U4 and U3 x86
- **Hypervisors:** The following hypervisors have been certified in 6.0.1.
 - VMware ESX 6.5 U1, U2, U3
 - VMware ESX 6.7 U3

Release 6.0.0.0

Google Cloud Support: Delphix now supports running in Google Cloud Platform for existing supported databases.

Enhanced Networking Adapter (ENA) Support: Delphix supports networking on AWS instances with the Elastic Network Adapter (ENA). This offers our customers enhanced networking capabilities and more economical options. Notably, this includes the AWS R4 instance types.

- **Masking NFS/CIFS Mount:** Our customers increasingly are masking files alongside their databases. The masking engine has classically supported this via FTP/SFTP but now to make things easier Delphix has introduced the ability to directly mount and mask a file system - over NFS and CIFS. This should dramatically simplify the process of file masking.
- **Oracle Quality:** Continued focus on Oracle quality and have introduced several quality improvements with our 6.0 release.

- **Masking API Updates:** 6.0 introduces a significant number of new endpoints, including mainframe control, as well as updates for existing endpoints. This release also introduces versioning for the masking API, allowing our customers to upgrade without risk of breaking their integrations.
- **AdoptOpenJDK 8 for the Delphix Toolkit:** Delphix has changed the Java Development Kit (JDK) that is included with the toolkit, and is sent to all Delphix connected environments. Customers who require using Oracle Java may continue to do so with the feature to provide their own Java, which shipped in 5.3.5.
- **Removed Instance Check:** When running in AWS or Azure, the product will no longer raise a fault when it detects that it is running on an unsupported instance. This enables Delphix to certify previously released software on new instances without having to modify the software.

The product will still detect what instance it is running on and include this information in the user interface and phone home bundles. We will also continue to publish a matrix of supported instances for Azure, AWS, and GCP in the product documentation. Delphix provides no guarantee of performance or support for unsupported instance types.
- **Upgrade Process:** The upgrade to 6.0 will be an in-place upgrade like other Delphix releases, there are a few changes that will improve the process overall for 6.0:
 - We will require an upgrade to an interim release first (either 5.3.6). This can be done at the same time as the customer upgrades to 6.0 or in the months prior.
 - We will be introducing new upgrade checks to ensure that customers are not using features that have been removed. For a list of removed features see [Deprecated and End-of-Life Features](#).
 - We will provide an upgrade image specific for each platform we support with Virtualization (VMWare, AWS, Azure, GCP). This will allow us to be more precise in customizing the images for each.

API Changes

This section covers the following topics:

- [API Changes in Delphix 6.0.13.0](#)
- [API Changes in Delphix 6.0.12.0](#)
- [API Changes in Delphix 6.0.11.0](#)
- [API Changes in Delphix 6.0.10.0](#)
- [API Changes in Delphix 6.0.9.0](#)
- [API Changes in Delphix 6.0.8.0](#)
- [API Changes in Delphix 6.0.7.0](#)
- [API Changes in Delphix 6.0.6.0](#)
- [API Changes in Delphix 6.0.5.0](#)
- [API Changes in Delphix 6.0.4.0](#)
- [API Changes in Delphix 6.0.3.0](#)

API Changes in Delphix 6.0.13.0

In Delphix 6.0.13.0, the new API version is 1.11.13. This section describes all API changes since API version 1.11.12, which was released with Delphix 6.0.12.0. All URL paths are relative to `/resources/json/delphix`.

What's Changed?

API Object	Path	Type	Name	Change
MSSqlNoBackupSyncParameters	N/A (value type)	API Type	MSSqlNoBackupSyncParameters	STAGING_PUSH_INGESTION feature flag removed
MSSqlSourcelessSyncStrategy	N/A (value type)	API Type	MSSqlSourcelessSyncStrategy	STAGING_PUSH_INGESTION feature flag removed
MSSqlStagingPushSyncStrategy	N/A (value type)	API Type	MSSqlStagingPushSyncStrategy	STAGING_PUSH_INGESTION feature flag removed
NetworkInterface	/network/interface	API Type	NetworkInterface	Removed Domain User access on CLI
NetworkRoute	/network/route	API Type	NetworkRoute	Removed Domain User access on CLI
NetworkRoute	/network/route	Property	protocol	Changed from “readonly” to “optional” when creating.
OAuth2Config	/service/oauth2	API Type	OAuth2Config	Added Domain User access from CLI. Provide System user only access for updates and root operations.

API Object	Path	Type	Name	Change
ObjectStore	N/A (value type)	API Type	ObjectStore	DELPHIX_DATA_BANK feature flag removed. Added update operation.
OracleDBConfig	N/A (value type)	Properties	databaseName uniqueName	Replace property pattern with format - oracleDatabaseName oracleDbUniqueName
OracleInstance	N/A (value type)	Property	instanceName	Replace property pattern with format - oracleInstanceName
OraclePDBConfig	N/A (value type)	Property	databaseName	Replace property pattern with format - oraclePDBName
S3ObjectStoreAccessInstanceProfile	N/A (value type)	API Type	S3ObjectStoreAccessInstanceProfile	DELPHIX_DATA_BANK feature flag removed.
S3ObjectStoreAccessKey	N/A (value type)	API Type	S3ObjectStoreAccessKey	DELPHIX_DATA_BANK feature flag removed.
S3ObjectStoreAccess	N/A (value type)	API Type	S3ObjectStoreAccess	DELPHIX_DATA_BANK feature flag removed.
S3ObjectStoreTestResult	N/A (value type)	API Type	S3ObjectStoreTestResult	DELPHIX_DATA_BANK feature flag removed.
S3ObjectStoreTest	N/A (value type)	API Type	S3ObjectStoreTest	DELPHIX_DATA_BANK feature flag removed.

API Object	Path	Type	Name	Change
S3objectStore	/storage/objectStorage	API Type	S3objectStore	DELPHIX_DATA_BANK feature flag removed. “Update” operation added for all the properties.

What's New?

API Object	Path	Type	Name	Description
FluentdAttributeDefinition	N/A (value type)	API Type	FluentdAttributeDefinition	Fluentd attribute definition.
FluentdAttribute	N/A (value type)	API Type	FluentdAttribute	Fluentd attribute.
FluentdConfig	/service/fluentd/configuration	API Type	FluentdConfig	Fluentd configuration information.
FluentdPlugin	/service/fluentd/plugins	API Type	FluentdPlugin	Upload and manage fluentd plugins.
FluentdRegularAttribute	N/A (value type)	API Type	FluentdRegularAttribute	Fluentd attribute with a plain value.
FluentdSecretAttribute	N/A (value type)	API Type	FluentdSecretAttribute	Fluentd attribute with a secret value.

API Object	Path	Type	Name	Description
MaskingServiceConfig	/maskingjob/serviceconfig	Property	maxJobFetchCount	Maximum number of jobs to fetch from masking service. Defaults to 500.
NettyVersionInfo	/netty	API Type	NettyVersionInfo	View netty version and switch between legacy and latest netty version.
ObjectStore	N/A (value type)	Property	configured	States whether an object store has been configured.
OracleBaseStagingLinkData	N/A (value type)	API Type	OracleBaseStagingLinkData	Represents common parameters to link an Oracle database using a staging database.
OracleLinkFromStaging	N/A (value type)	API Type	OracleLinkFromStaging	Represents parameters to link a non-pluggable Oracle database using a staging database.
OraclePDBLinkFromStaging	N/A (value type)	API Type	OraclePDBLinkFromStaging	Represents parameters to link a pluggable Oracle database using a staging database.
OracleSourceLessSyncStrategy	N/A (value type)	API Type	OracleSourceLessSyncStrategy	Base type for Oracle source less sync strategy and associated parameters.
OracleStagingPushSyncParameters	N/A (value type)	API Type	OracleStagingPushSyncParameters	The parameters to use as input to sync a staging Oracle database.

API Object	Path	Type	Name	Description
OracleStagingPushSyncStrategy	N/A (value type)	API Type	OracleStagingPushSyncStrategy	Oracle specific parameters for staging push sync strategy.
OracleStagingSourceParameters	N/A (value type)	Property	instanceName	The name (sid) of the instance.
S3ObjectStore	/storage/objectStorage	Root operations	cacheHitsReport clearCacheHits	Get print out of the ZettacCache hit-by-size histogram. Clear the current ZettacCache hit-by-size histogram.
S3ObjectStoreTestResult	N/A (value type)	Property	errorMessage	Error message from connectivity test.

API Changes in Delphix 6.0.12.0

In Delphix 6.0.12.0, the new API version is 1.11.12. This section describes all API changes since API version 1.11.11, which was released with Delphix 6.0.11.0.

All URL paths are relative to `/resources/json/delphix`.

What's Changed?

API Object	Path	Type	Name	Change
<code>AbstractToolkit</code>	<code>/toolkit</code>	property	<code>language</code>	Removed.
<code>CipherSuite</code>	<code>/service/tls/cipherSuite</code>	property	<code>name</code>	Removed values for enum - <code>"TLS_RSA_WITH_AES_256_CBC_SHA"</code> , <code>"TLS_RSA_WITH_AES_128_CBC_SHA"</code> , <code>"TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA"</code> , <code>"TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA"</code> , <code>"TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA"</code> , <code>"TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA"</code> , <code>"TLS_DHE_RSA_WITH_AES_256_CBC_SHA"</code> , <code>"TLS_DHE_RSA_WITH_AES_128_CBC_SHA"</code> , <code>"TLS_DHE_DSS_WITH_AES_256_CBC_SHA"</code> , <code>"TLS_DHE_DSS_WITH_AES_128_CBC_SHA"</code> .
<code>HttpConnectorConfig</code>	<code>/service/httpConnector</code>	property	<code>tlsVersions</code>	Removed values for enum - <code>"SSLv3"</code> , <code>"TLSv1"</code> , <code>"TLSv1_1"</code>

API Object	Path	Type	Name	Change
OAuth2Config	/service/ oauth2	property	userMatchingFieldType	Updated description.
OracleLinkedSource	N/A (value type)	property	config backupLevelEnabled rmanChannels filesPerSet checkLogical encryptedLinkingEnabled compressedLinkingEnabled bandwidthLimit numberOfConnections	Removed.
OracleSource	N/A (value type)	property	config	Removed.

API Object	Path	Type	Name	Change
StorageDevice	/storage/ device	property	list operations	Removed CliDisplay parameters - "bootDevice", "allocating". Removed operations - "remove", "removeVerify".
UnixHost	N/A (value type)	property	oracleJdbc KeystorePas sword	Removed.
User	/user	property	apiUser	Renamed to allowPasswordAuthentication (see in What's New?)

What's New?

API Object	Path	Type	Name	Description
Configured StorageDev ice	N/A (value type)	prop erty	allocating	True if the device is being used for allocations. Allocations will be disabled automatically if the device is going to be removed. The allocating property will also be false for devices that are not configured.
MSSqlNoBac kupSyncPar ameters	N/A (value type)	API Typ e	MSSqlNoBacku pSyncParamet ers	The parameters to use as input to sync MSSQL databases without a backup.
MSSqlSourc elessSyncS trategy	N/A (value type)	API Typ e	MSSqlSourceL essSyncStrat egy	MSSQL specific parameters for sourceless sync strategy.

API Object	Path	Type	Name	Description
MSSqlStagingPushSyncStrategy	N/A (value type)	API Type	MSSqlStagingPushSyncStrategy	MSSQL specific parameters for staging push sync strategy.
OracleHostParameters	N/A (value type)	API Type	OracleHostParameters	Oracle specific host parameters.
OracleLinkedSource	N/A (value type)	property	syncStrategy overrideMapsTo	Parameters used to sync the container. These are persisted and used for every sync operation. Override 'mapsTo' property from source.js for the query parameters defined in 'list' section.
OracleManagedSource	N/A (value type)	property	config	Reference to the configuration for the source.
Plugin	N/A (value type)	property	language	Implementation language for workflows in this plugin.
StorageDeviceRemovalStatus	/storage/remove	property	rootOperations	Add "start" and "verify" root operations to <ul style="list-style-type: none"> remove storage devices verify storage devices can be removed.
StorageDeviceRemoveParameters	N/A (value type)	API Type	StorageDeviceRemoveParameters	The parameters to use as input when removing devices.
UnixHost	N/A (value type)	property	oracleHostParameters	The Oracle specific parameters associated with the host.
User	/user	property	allowPasswordAuthentication	If the Delphix Engine has been configured for SAML SSO, only users with this property set can authenticate with a username and password for API or CLI access.

API Changes in Delphix 6.0.11.0

In Delphix 6.0.11.0, the new API version is 1.11.11. This section describes all API changes since API version 1.11.10, which was released with Delphix 6.0.10.0.

All URL paths are relative to `/resources/json/delphix`.

What's Changed?

API Object	Path	Type	Name	Change
AdvancedSettingsInfo	/system/advancedSettings	API Type	AdvancedSettingsInfo	Hidden behind feature flag <code>SYSTEMTUNABLEINTERFACE</code>
DelphixManagedBackupIngestionStrategy	N/A (value type)	API Type	DelphixManagedBackupIngestionStrategy	Removed.
ExternalBackupIngestionStrategy	N/A (value type)	API Type	ExternalBackupIngestionStrategy	Removed.
IngestionStrategy	N/A (value type)	API Type	IngestionStrategy	Removed.
MSSqlDBConfig	N/A (value type)	property	mssqlUser	Removed.

API Object	Path	Type	Name	Change
MSSqlLinkedSource	N/A (value type)	property	config sharedBackupLocations ingestionStrategy mssqlNetbackupConfig mssqlDefaultConfig	Removed.
MSSqlSourceSyncStrategy	N/A (value type)	property	config	Optional while creating.
NoBackupIngestionStrategy	N/A (value type)	API Type	NoBackupIngestionStrategy	Removed.
OracleLinkFromExternal	N/A (value type)	property	config	Extends reference changed from OracleBaseLinkData to OracleBaseExternalLinkData. Config property removed.
OraclePDBLinkFromExternal	N/A (value type)	property	config	Extends reference changed from OracleBaseLinkData to OracleBaseExternalLinkData. Config property removed.

API Object	Path	Type	Name	Change
OracleTimeflow	N/A (value type)	property	tdeUUID	No Longer behind feature flag ORACLEMTTDE .
OracleVirtualPdb Source	N/A (value type)	property	parentTdeKeystorePath parentTdeKeystorePassword tdeExportedKeyFileSecret tdeUUID	No Longer behind feature flag ORACLEMTTDE .
Source	N/A (value type)	property	config	nameParent changed from config to container . Config property removed from properties and list operation.
SsoFailedLoginRecord	N/A (value type)	property	userAgent originIp reason	Extends reference changed from TypedObject to FailedLoginRecord . All properties removed.
SsoSuccessfulLoginRecord	N/A (value type)	property	userAgent originIp	Extends reference changed from TypedObject to LoginRecord . Properties removed.

What's New?

API Object	Path	Type	Name	Description
AppDataCompatibilityCriteria	N/A (value type)	API Type	AppDataCompatibilityCriteria	The compatibility criteria to use for filtering the list of available AppData repositories.
AppDataSource	N/A (value type)	property	config	Reference to the configuration for the source.
ASESource	N/A (value type)	property	config	Reference to the configuration for the source.
FailedLoginRecord	N/A (value type)	API Type	FailedLoginRecord	Represents an successful login using an external identity provider.
LoginRecord	N/A (value type)	API Type	LoginRecord	Represents a failed SSO login.
MSSqlDelphixManagedSyncStrategy	N/A (value type)	property	config	Reference to the configuration for the source.
MSSqlExternalManagedSourceSyncStrategy	N/A (value type)	property	config	Reference to the configuration for the source.
MSSqlLinkedSource	N/A (value type)	property	syncStrategy overrideMapsTo	Sync strategy for this source. Override mapsTo property for the query parameter defined in list at the root type.
MSSqlStagingSource	N/A (value type)	property	config	Reference to the configuration for the source.

API Object	Path	Type	Name	Description
MSSqlUser	N/A (value type)	property	config	Reference to the configuration for the source.
MySQLSource	N/A (value type)	property	config	Reference to the configuration for the source.
OAuth2Config	/service/oauth2	API Type	OAuth2Config	OAuth2 Configuration.
OAuth2FailedLoginRecord	N/A (value type)	API Type	OAuth2FailedLoginRecord	Represents a failed OAuth2 login.
OAuth2SuccessfulLoginRecord	N/A (value type)	API Type	OAuth2SuccessfulLoginRecord	Represents a successful OAuth2 login.
OracleBaseExternalLinkData	N/A (value type)	API Type	OracleBaseExternalLinkData	Represents common parameters to link all externally managed Oracle databases.
OracleSourceBasedSyncStrategy	N/A (value type)	API Type	OracleSourceBasedSyncStrategy	Base type for Oracle source based sync strategy and associated parameters.
OracleSource	N/A (value type)	property	config	Reference to the configuration for the source.
OracleSyncStrategy	N/A (value type)	API Type	OracleSyncStrategy	Base type for Oracle sync strategy and associated parameters.

API Object	Path	Type	Name	Description
OracleVirtualPdbSource	N/A (value type)	property	tdeKeyIdentifier	ID of the key created by Delphix, as recorded in encryption_keys.key_id.
OracleVirtualSource	N/A (value type)	property	newDBID	Indicates whether Delphix will generate a new DBID during VDB provision or refresh.
OsAdminStatus	/osadmin/engineStatus	API Type	OsAdminStatus	Information for the current state of the Delphix Engine.
PgSQLSource	N/A (value type)	property	config	Reference to the configuration for the source.
PluginManifest	N/A (value type)	property	hasVirtualCleanup	Indicates whether virtual.cleanup() operation has been implemented.
ReplicationSourceState	N/A (value type)	property	lastKnownTargetVersion	The Delphix version of the target engine during the last replication using the spec.
ReplicationTargetState	N/A (value type)	property	lastKnownSourceVersion	The Delphix version of the source engine during the last replication of the namespace.

API Changes in Delphix 6.0.10.0

In Delphix 6.0.10.0, the new API version is 1.11.10. This section describes all API changes since API version 1.11.9, which was released with Delphix 6.0.9.0.

All URL paths are relative to `/resources/json/delphix`.

What's Changed?

API Object	Path	Type	Name	Change
<code>CredentialsEnvVars</code>	N/A (value type)	API Type	<code>CredentialsEnvVars</code>	Extends reference changed from <code>TypedObject</code> to <code>IdentifiableArrayElement</code> .
<code>CredentialsEnvVars</code>	N/A (value type)	property	<code>baseVarName</code> <code>credentials</code>	Optional for an update. Optional for update.
<code>IscsiTarget</code>	<code>/storage/iscsi/target</code>	property	<code>state</code>	Added enum value <code>ERROR</code> .
<code>KeyPairCredential</code>	N/A (value type)	property	<code>privateKey</code> <code>publicKey</code>	Required for create. Optional for update.
<code>NetworkRoute</code>	<code>/network/route</code>	property	<code>list</code>	Add enum value <code>protocol</code> in " <code>cliDisplay</code> ".
<code>Operation</code>	N/A (value type)	API Type	<code>Operation</code>	Extends reference changed from <code>TypedObject</code> to <code>IdentifiableArrayElement</code> .

API Object	Path	Type	Name	Change
OracleDatabaseContainer	N/A (value type)	property	racMaxInstanceLag	FeatureFlag , ORACLERACMAXINSTANCEL AG ' is removed. default value changed to 3.
OracleExportParameters	N/A (value type)	property	filesystemLayout	Extends reference changed from TimeFlowFilesystemLayout to OracleExportTimeflowF ilesystemLayout
SystemInitializationParameters	N/A (value type)	property	devices	Optional while creating.
TimeZone	/timezone	property	id	Removed values for enum - "ACT", "AET", "AGT", "ART", "AST", "SST", "US/Pacific -New", "VST".

What's New?

API Object	Path	Type	Name	Description
IdentifiableArrayElement	N/A (value type)	API Type	IdentifiableArrayElement	A unique identifier is generated by the engine when this object is in an array. If this value is set and this object is in an update request, this id identifies the element to update in the existing array. In that case, this object can be sparse (only the changed properties need to be included). The resulting position of the element in the array is its position in the update. If an element with this id does not exist, this object is ignored. If the id is not set, a new object is added to the array and a new id is assigned to it by the engine. If an existing element id is omitted from an update, the corresponding element is deleted. To preserve an existing array element unchanged, the API client must include an element in the update with just this id.
MSSqlSnapshot	N/A (value type)	property	backupSoftwareType backupLocationType	The backup software was used to restore the source database backup for this snapshot. The type of backup location where the source database backup was present.
NetworkRoute	/network/route	property	protocol create return	The protocol is used to manage the route. Create a new NetworkRoute object. Return the NetworkRoute object reference.
ObjectStore	N/A (value type)	API Type	ObjectStore	An object store.
OracleExportTimeflowFilesystemLayout	N/A (value type)	API Type	OracleExportTimeflowFilesystemLayout	A filesystem layout that matches the filesystem of a Delphix Oracle TimeFlow for use in database export.

API Object	Path	Type	Name	Description
S3objectStoreAccessInstanceProfile	N/A (value type)	API Type	S3objectStoreAccessInstanceProfile	S3 object store access through instance profile.
S3objectStoreAccessKey	N/A (value type)	API Type	S3objectStoreAccessKey	S3 object store access key.
S3objectStoreAccess	N/A (value type)	API Type	S3objectStoreAccess	S3 Object-store access.
S3objectStoreTestResult	N/A (value type)	API Type	S3objectStoreTestResult	An Amazon S3 object store connectivity test result.
S3objectStoreTest	N/A (value type)	API Type	S3objectStoreTest	An Amazon S3 object store connectivity test object.
S3objectStore	N/A (value type)	API Type	S3objectStore	An Amazon Simple Storage Service (Amazon S3) object store.
SystemInitializationParameters	N/A (value type)	property	objectStore	Object storage.

API Changes in Delphix 6.0.9.0

In Delphix 6.0.9.0, the new API version is 1.11.9. This section describes all the API changes since API version 1.11.8, which was released with Delphix 6.0.8.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
AppDataProvisionParameters	N/A (value type)	property	sourceConfig	Type changed from <code>AppDataSourceConfig</code> to <code>AppDataDirectSourceConfig</code> .
BaseSupportBundleParameters	N/A (value type)	property	bundleType	Added enum value <code>DOCKER_LOG</code> .
CipherSuite	<code>/service/tls/cipherSuite</code>	property	name	Added enum values <code>TLS_AES_128_CCM_8_SHA256</code> , <code>TLS_AES_128_CCM_SHA256</code> , <code>TLS_AES_128_GCM_SHA256</code> , <code>TLS_AES_256_GCM_SHA384</code> , and <code>TLS_CHACHA20_POLY1305_SHA256</code> .
HttpConnectorConfig	<code>/service/httpConnector</code>	property	tlsVersions	Added enum value <code>TLSv1_3</code> .
Namespace	<code>/namespace</code>	property	locked	No longer behind a feature flag.
OracleCluster	<code>/environment</code>	property	name version	No longer read-only. Deleted.

API Object	Path	Type	Name	Change
Replication Spec	N/A (value type)	property	lockedProfile	No longer behind a feature flag.
WindowsCluster	/environment	property	name	No longer read-only.

What Is New

API Object	Path	Type	Name	Description
AlertProfile	/alert/profile	property	user	User to which the alert profile is assigned. Defaults to the logged-in user.
CredentialsEnvVars	N/A (value type)	API type	CredentialsEnvVars	Credentials to be placed in environment variables for an operation.
IscsiTargetDiscoveryParameters	N/A (value type)	property	chapUsername chapPassword chapUsernameIn chapPasswordIn	CHAP username to be used for iSCSI Discovery. CHAP password to be used for iSCSI Discovery. Target/Mutual CHAP username to be used for iSCSI Discovery (bidirectional authentication). Target/Mutual CHAP password to be used for iSCSI Discovery (bidirectional authentication).

API Object	Path	Type	Name	Description
IscsiTarget	/storage/ iscsi/ target	property	chapUsername chapPassword chapUsernameIn chapPasswordIn	CHAP username to be used for iSCSI Target authentication. CHAP password to be used for iSCSI Target authentication. Target/Mutual CHAP username (bidirectional authentication). Target/Mutual CHAP password (bidirectional authentication).
SourceOperation	N/A (value type)	property	credentials EnvVarsList	List of environment variables that will contain credentials for this operation.
SystemVersion	/system/ version	property	hotfixVersion	The hotfix version.

API Changes in Delphix 6.0.8.0

In Delphix 6.0.8.0, the new API version is 1.11.8. This section describes all API changes since API version 1.11.7, which was released with Delphix 6.0.7.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
<code>AttachData</code>	N/A (value type)	property	<code>config</code>	Deleted.
<code>Container</code>	<code>/database</code>	operation	<code>undo</code>	HTTP method changed from <code>GET</code> to <code>POST</code> .
<code>LinkData</code>	N/A (value type)	property	<code>config</code>	Deleted.
<code>MSSqlAttachData</code>	N/A (value type)	property	<code>backupLocationCredentials</code> <code>backupLocationUser/source</code> <code>config</code> <code>externalFilePath</code> <code>ingestionStrategy</code> <code>mssqlCommvaultConfig</code> <code>mssqlNetbackupConfig</code> <code>mssqlUser</code> <code>sharedBackupLocations</code>	Deleted. (See the new <code>syncStrategy</code> property in the table below.)

API Object	Path	Type	Name	Change
MSSqLLinkData	N/A (value type)	property	backupLocationC redentials backupLocationU ser config externalFilePat h ingestionStrate gy mssqlCommvaultC onfig mssqlNetbackupC onfig mssqlUser sharedBackupLoc ations	Deleted. (See the new syncStrategy property in the table below.)
MSSqLLinkedSource	/source	property	backupLocationC redentials backupLocationU ser externalFilePat h	Deleted.
NamespaceFailoverParameters	N/A (value type)	property	smartFailover	Default changed to true.
OracleSyncFromExternalParameters	N/A (value type)	property	filesForFullBackup	No longer behind a feature flag.
SourceIngestionData	N/A (value type)	property	containerType	Removed obsolete value VMWARE_CONTAINER from enum.

API Object	Path	Type	Name	Change
SourceRepository	/repository	property	LinkingEnabled	Deleted.
SourceTypeAggregateIngestedSize	N/A (value type)	property	containerType	Removed obsolete value <code>VMWARE_CONTAINER</code> from enum.
Source	/source	property	status	Removed value <code>DETACHED</code> from enum.
VMware*		API types	VMware*	Deleted all types whose names start with <code>VMware</code> (obsolete).

What Is New

API Object	Path	Type	Name	Description
AdvancedSettingsInfo	/system/advancedSettings	API type	AdvancedSettingsInfo	Set advanced virtualization, OS, network, and service tunables.
ASESnapshot	/snapshot	property	dbEncryptionKey	Name of database encryption key present for this snapshot.
IscsiInitiator	/storage/iscsi/initiator	API type	IscsiInitiator	Endpoint for iSCSI initiator information.

API Object	Path	Type	Name	Description
<code>MSSqlAttachData</code>	N/A (value type)	property	<code>syncStrategy</code>	Configuration that determines what sync strategy the source will use.
<code>MSSqlDelphixManagedSyncStrategy</code>	N/A (value type)	API type	<code>MSSqlDelphixManagedSyncStrategy</code>	MSSQL specific parameters for delphix managed source based sync strategy.
<code>MSSqlExternalManagedSourceSyncStrategy</code>	N/A (value type)	API type	<code>MSSqlExternalManagedSourceSyncStrategy</code>	MSSQL specific parameters for externally managed source based sync strategy.
<code>MSSqlLinkData</code>	N/A (value type)	property	<code>syncStrategy</code>	Configuration that determines what sync strategy the source will use for linking.
<code>MSSqlSourceSyncStrategy</code>	N/A (value type)	API type	<code>MSSqlSourceSyncStrategy</code>	MSSQL specific parameters for source based sync strategy.
<code>MSSqlSyncStrategy</code>	N/A (value type)	API type	<code>MSSqlSyncStrategy</code>	Base type for mssql specific parameters for sync strategy.
<code>Namespace</code>	<code>/ namespace</code>	property	<code>locked</code>	Indicates the namespace is locked. (Behind the <code>DATA_VAULT</code> feature flag.)
<code>OracleDatabaseContainer</code>	<code>/ container</code>	property	<code>racMaxInstanceLag</code>	Maximum number of log sequences that a node can lag behind on RAC before considering instance offline. (Behind the <code>ORACLERACMAXINSTANCELAG</code> feature flag.)
<code>ReplicationSpec</code>	N/A (value type)	property	<code>lockedProfile</code>	Indicates the replication profile is locked. (Behind the <code>DATA_VAULT</code> feature flag.)

API Object	Path	Type	Name	Description
SyncStrategy	N/A (value type)	API type	SyncStrategy	Base type for specific parameters for sync strategy.
Tunable	N/A (value type)	API type	Tunable	The name of the tunable and its value.
TunableIdentifier	N/A (value type)	API type	TunableIdentifier	The subsystem and name for a tunable.

API Changes in Delphix 6.0.7.0

In Delphix 6.0.7.0, the new API version is 1.11.7. This section describes all API changes since API version 1.11.6, which was released with Delphix 6.0.6.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
PgSQLDBClusterConfigConnectivity	N/A (value type)	API type	PgSQLDBClusterConfigConnectivity	Deleted.
OracleBaseDBConfig	<code>/sourceconfig</code>	property	nonSysCredentials	Changed type from PasswordCredential to Credential.
OracleLinkFromExternal	N/A (value type)	property	nonSysCredentials	Changed type from PasswordCredential to Credential.
ReplicationSecureList	N/A (value type)	API type	ReplicationSecureList	No longer behind feature flag MDD.
SNMPV3Manager	<code>/service/snmp/v3/manager</code>	property	username	Introduced a minimum length of 1.
Source	<code>/source</code>	property	status	Added DETACHED to the list of possible values.

What Is New

API Object	Path	Type	Name	Description
AbstractToolkit	/ toolkit	root operation	schemaDefinitions	Get the platform's JSON schema definitions that plugin schemas can reference.
ASESyncParameters	N/A (value type)	property	dropAndRecreateDevices	If this parameter is set to true, it will drop the older devices and create new devices during manual sync operations instead of trying to remap the devices. This might increase the space utilization on Delphix Engine.
NamedKeyPairCredential	N/A (value type)	API type	NamedKeyPairCredential	Username and key-pair credential.
NamedPasswordCredential	N/A (value type)	API type	NamedPasswordCredential	Pair of username and password credential.
OracleSyncFromExternalParameters	N/A (value type)	property	filesForFullBackup	List of datafiles to take a full backup of. This would be useful in situations where certain datafiles could not be backed up during previous SnapSync due to corruption or because they went offline.
SsoConfig	/ service /sso	property	entityId	Audience Restriction (SP entity ID, Partner's Entity ID) of this engine as an SSO service provider.
SsoConfig	/ service /sso	property	responseSkewTime	Maximum time difference allowed between a SAML response and the engine's current time, in seconds. If not set, it defaults to 86,400 seconds (one day).

API Object	Path	Type	Name	Description
SsoConfiguration	/service/sso	property	maxAuthenticationAge	How far in the past to accept authentications to the identity provider, in seconds. If not set, it defaults to 120 seconds.
SupportBundleConfiguration	/service/support/bundle	property	maxActions	Maximum number of actions not referenced anywhere else to include from the metadata store. The most recent such actions are included.
SupportBundleConfiguration	/service/support/bundle	operation	read	Retrieve the specified SupportBundleConfiguration object.
SupportBundleConfiguration	/service/support/bundle	operation	update	Update the specified SupportBundleConfiguration object.

API Changes in Delphix 6.0.6.0

In Delphix 6.0.6.0, the new API version is 1.11.6. This section describes all API changes since API version 1.11.5, which was released with Delphix 6.0.5.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
OracleMultitenantProvisionParameters	N/A (value object)	object property	timeflowPointParameters	The source timeflow point can now reference an Oracle non-multitenant database. In which case the non-multitenant database will be provisioned from the timeflow point to a virtual pluggable database.

API Changes in Delphix 6.0.5.0

In Delphix 6.0.5.0, the new API version is 1.11.5. This section describes all API changes since API version 1.11.4, which was released with Delphix 6.0.4.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
CloudStatus	/service/ cloud	property	delphixDataServicesComponentStatus	Made read-only for update.
CloudStatus	/service/ cloud	property	delphixDataServicesComponentInfo	Made read-only for update.
CloudStatus	/service/ cloud	root operation	enable	Added optional payload <code>CloudEnableParameters</code> .
OracleDBConfig	/ sourceconfig	property	tdeKeystorePassword	Impose a maximum length of 256.
OracleInstance	N/A (value object)	property	instanceNumber	Change type from number to integer.
OracleSnapshot	/snapshot	property	redoLogSizeInBytes	Change type from number to integer.
OracleStartParameters	N/A (value object)	property	instances	Change type from number to integer.

API Object	Path	Type	Name	Change
OracleStopParameters	N/A (value object)	property	instances	Change type from number to integer.
TimeZone	/timezone	property	id	Added enum values America/Nuuk and Asia/Qostanay.

What Is New

API Object	Path	Type	Name	Description
CloudEnableParameters	N/A (value object)	value type	CloudEnableParameters	Parameters to the Cloud Enable operation
CloudStatus	/service/cloud	property	proxyMode	Whether an HTTP proxy must be used to connect to Central Management.
CloudStatus	/service/cloud	property	proxyConfiguration	Proxy configuration for communication with Delphix Central Management. This property is ignored unless the 'proxyMode' property is set to CLOUD_SPECIFIC_SETTING.
CloudStatus	/service/cloud	operation	update	Update the specified CloudStatus object. Payload is a CloudStatus.
HostConfiguration	N/A (value object)	property	powerShellVersion	The PowerShell version installed on the windows target host.
OracleSTConvertedToPDBAttachData	N/A (value object)	value type	OracleSTConvertedToPDBAttachData	Sub-type of OraclePDBAttachData.

API Object	Path	Type	Name	Description
OracleTimeflow	/timeflow	property	tdeUUID	Unique identifier for TimeFlow-specific TDE objects that reside outside of Delphix storage.
OracleVirtualPdbSource	/source	property	parentTdeKeystorePassword	The password of the keystore specified in parentTdeKeystorePath.
OracleVirtualPdbSource	/source	property	tdeExportedKeyFileSecret	Secret to be used while exporting and importing vPDB encryption keys if Transparent Data Encryption is enabled on the vPDB.
OracleVirtualPdbSource	/source	property	tdeUUID	Unique identifier for PDB-specific TDE objects that reside outside of Delphix storage.
StaticHostAddress	/service/host/address	new API	StaticHostAddress	Static mapping of hostname to IP address.
WindowsHost	/host	property	connectorVersion	The Windows Connector version that is installed on the provided host.
WindowsHost	/host	property	connectorDotNetFrameworkVersion	The .NET Framework version used for Windows Connector Service.

API Changes in Delphix 6.0.4.0

In Delphix 6.0.4.0, the new API version is 1.11.4. This section describes all API changes since API version 1.11.3, which was released with Delphix 6.0.3.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
OracleDBConfig Connectivity	N/A (value object)	object property	password	Name changed to credentials Type changed to Credential
OracleBaseAtta chData	N/A (value object)	object property	oracleFallback Credentials	Type changed to Credential
OracleBaseDBCo nfig	/ sourceconfig	object property	credentials	Type changed to Credential
OracleBaseLink Data	N/A (value object)	object property	oracleFallback Credentials	Type changed to Credential
OracleDatabase Container	/database	object property	tdeProvisionin gEnabled	Deleted
OracleManagedS ource	/source	object property	mountBase	Format is now unixrestricted path
OracleVirtualP dbSource	/source	object property	parentTdeKeyst orePath	Format is now unixrestricted path
UnixRuntimeMou ntInformation	N/A (value object)	object property	nfsVersionReas on	Added enum value INCOMPLETE_V4_ CONFIG

API Object	Path	Type	Name	Change
X500DistinguishedNameComposite	N/A (value object)	object property	dnname	Minimum length is now 1

What Is New

API Object	Path	Type	Name	Description
X509Certificate	/service/certificate	new API	X509Certificate	X509 Certificate.

API Changes in Delphix 6.0.3.0

In Delphix 6.0.3.0, the new API version is 1.11.3. This section describes all API changes since API version 1.11.2, which was released with Delphix 6.0.2.0.

All URL paths are relative to `/resources/json/delphix`.

What Changed

API Object	Path	Type	Name	Change
ApplyVersionParameters	N/A (value object)	object property	verify	Default changed from true to false.
OracleBaseAttachData	N/A (value object)	object property	dbUser dbCredentials	Name changed to oracleFallbackUser. The database user. Optional if bequeath connections are enabled (to be used in case of bequeath connection failures). Name changed to oracleFallbackCredentials. The password for the database user. Optional if bequeath connections are enabled (to be used in case of bequeath connection failures).
OracleBaseLinkData	N/A (value object)	object property	dbUser dbCredentials	Name changed to oracleFallbackUser. The database user. Optional if bequeath connections are enabled (to be used in case of bequeath connection failures). Name changed to oracleFallbackCredentials. The password for the database user. Optional if bequeath connections are enabled (to be used in case of bequeath connection failures).
Plugin	/toolkit	object property	status	Deleted.

What Is New

API Object	Path	Type	Name	Description
AbstractToolkit	/toolkit	object property	status	The status of the toolkit. ACTIVE indicates toolkit is actively referenced and in use. INACTIVE means toolkit needs to go through a manual upgrade operation before it can be used.
Certificate	/service/tls/caCertificate /service/tls/endEntityCertificate	object property	isCertificateAuthority	Whether this certificate is a Certificate Authority (CA).
ConfiguredStorageDevice	/storage/device	object property	fragmentation	Percent fragmentation for this device.
CyberArkVaultCredential	N/A (value object)	subclass of Credential	CyberArkVaultCredential	CyberArk vault based security credential.
DNSConfig	/service/dns	object property	source	The source of the DNS configuration (STATIC or DHCP).
HashiCorpVaultCredential	N/A (value object)	subclass of Credential	HashiCorpVaultCredential	HashiCorp vault based security credential.
LicenseInfo	/license	object	LicenseInfo	Retrieve all external licenses.
NetworkDSPTestParameters	N/A (value object)	object property	xportScheduler	The transport scheduler to use.
OracleInstall	/repository	object property	oracleBaseConfig oracleBaseHome	The Oracle Base Config directory. The Oracle Base Home directory.

API Object	Path	Type	Name	Description
OracleVirtualPdbSource	/source	object property	parentTdeKeystorePath	Path to a copy of the parent's Oracle transparent data encryption keystore on the target host. Required to provision from snapshots containing encrypted database files.
PasswordVault	/service/passwordVault	new API	PasswordVault	Password vault configuration.
Plugin	/toolkit	object property	luaName minimumLuaVersion	The name of the LUA toolkit that this plugin can upgrade. The minimum version (in major.minor format) of a LUA toolkit that this plugin can upgrade.
RunDefaultPowerShellOnSourceOperation	N/A (value object)	subclass of SourceOperation	RunDefaultPowerShellOnSourceOperation	A user-specifiable operation that runs a PowerShell command (using default version) on the target host.
SnapshotCapacityData	/capacity/snapshot	object property list parameter	namespace snapshotTimezone snapshotLatestChange namespace	Reference to the namespace to which this snapshot belongs. Time zone of the source database at the time the snapshot was taken. The timestamp of the latestChangePoint of the associated snapshot. The namespace to list snapshot data for. If null, will limit returned snapshots to the default namespace.

API Object	Path	Type	Name	Description
StatisticSlice	/analytics	getData parameter	count	The number of data points to return. Mds data points will be combined in order to meet this requirement. When count is specified at least two of the other getData parameters must be specified.
SystemInfo	/system	object property	poolFragmentation	Percent fragmentation for the domain0 pool.

Support Matrices

This section covers the following topics:

- [Oracle Matrix](#)
- [SQL Server Matrix](#)
- [SAP ASE Matrix](#)
- [Db2 Matrix](#)
- [PostgreSQL Matrix](#)
- [HANA Matrix](#)
- [EBS Matrix](#)
- [Kerberos Support Matrix](#)
- [Data Source Certifications](#)
- [Virtualization Platform Matrix](#)
- [vFiles Matrix](#)
- [Windows Connector Matrix](#)

Oracle Matrix

Disclaimers and Support Notes

Delphix Support Policies specifically list **major** and **minor** release coverage. If a minor release is listed as covered, all patch releases under that minor release are certified. Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL), 64-bit OS support only.

The Oracle version in the support matrix applies to both Oracle DB and Oracle Grid. Delphix supports the same set of features and functionality for supported non-multitenant database versions. Currently, Delphix does not support the Oracle 12c feature of `THREADED_EXECUTION` being set to `TRUE`, because this disables OS authentication.

Summary of Delphix features that are unsupported for the Oracle12c and higher Multi-tenant configuration:

- Existing Virtual Container Database (vCDB) as target for provisioning an additional vPDB
- Customize VDB settings/initialization parameters. Includes the following:
 - Customize `init.ora` database parameter during provisioning
 - Config templates
 - Online redo log size
 - Number of RAC VDB instances
 - Online redo log groups
 - Archive log mode
 - Setting new DBID
 - Customize local listeners
- Virtual to Physical (V2P) Support
- Resumable initial SnapSync
- Validated Sync
- Source continuity for dSource upgraded from Oracle 12c non-multitenant to multitenant database
- Cross-platform provisioning (XPP) to the virtual database
- Oracle LiveSources

Delphix supports systems that are part of an Exadata or Exadata Cloud-at-Customer (ExaCC) cluster, provided that the operating system and DBMS version are in the supported list. Oracle Exadata, Exadata Cloud, and Exadata Cloud-at-Customer (ExaCC) offer several features that are reliant on the underlying Exadata storage. Delphix supports creating dSources from Exadata source clusters and provisioning VDBs and vPDBs to Exadata target clusters.

The following Exadata-only features will not be available in child VDBs running in Delphix storage:

- Smart Scan - allows SQL processing to be offloaded to Exadata storage cells. Queries that leverage Smart Scan will execute successfully in VDBs but query performance may be impacted.
- Smart Flash Cache - allows frequently used data blocks to be stored in the flash cache of Exadata storage cells to allow for faster access. Queries against objects that are normally cached in Smart Flash Cache will execute successfully in VDBs but query performance may be impacted. As the virtualization engine uses memory to cache frequently used data blocks, allocating extra memory to the engine may help to mitigate this performance impact.
- Hybrid Columnar Compression (HCC) - compresses data and stores it in a columnar format, allowing Exadata storage cells to deliver faster query performance while also saving on storage. Columns compressed using HCC cannot be used by queries executed in VDBs. Affected columns can be decompressed though care should be taken to evaluate the impact of doing so on VDB refresh times and Delphix storage usage. In Delphix VDBs, queries against partitions or tables using HCC will fail with the following error *"ORA-64307: Exadata Hybrid Columnar Compression is not supported for tablespaces on this storage type"*.

Color	Supported?
✓	Yes
✗	No
N/A	Not Applicable

Oracle Database Editions

Delphix supports the following Oracle database editions:

- Enterprise
- Standard

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please refer to [KBA 5622](#) for the required steps.

Supported OS Version	Supported DBMS Version						
	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 5.5	✓	✓	N/A	N/A	N/A	N/A	N/A
RHEL 5.6	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.7	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.8	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.9	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.10	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.11	✓	✓	✓	N/A	N/A	N/A	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 6.0	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.1	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.2	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.3	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.4	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.5	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.6	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.7	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.8	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.9	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.10	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.0	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.1	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.2	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.3	N/A	✓	✓	✓	✓	N/A	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 7.4	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.5	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.6	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.7	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.8	N/A	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	N/A
RHEL 7.9	N/A	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.5+	✓ 6.0.4+	N/A
RHEL 8.0	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.1	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.2	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.3	N/A	N/A	N/A	N/A	N/A	✓ 6.0.7+	✓ 6.0.11+
RHEL 8.4	N/A	N/A	N/A	N/A	N/A	✓ 6.0.10+	✓ 6.0.11+

SUSE Linux Enterprise Server (SLES)

 Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 11	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP1	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP2	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP3	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP4	✗	✓	✓	N/A	N/A	N/A	N/A
SLES 12	N/A	✓	✓	N/A	N/A	N/A	N/A
SLES 12 SP1	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP2	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP3	N/A	✓	✓	✓	✓	✓	N/A
SLES 12 SP4	N/A	✓	✓	✓	✓	✓	N/A
SLES 12 SP5	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+	✗	✓ 6.0.4+	N/A
SLES 15	N/A	✗	✗	✓	✓	✓	N/A
SLES 15 SP1	N/A	✗	✗	✗	✗	✓	✓ 6.0.11 +
SLES 15 SP2	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11 +

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 15 SP3	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11+

Solaris SPARC

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	✗	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i ¹ - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 SPARC kernel patch 144500, or a newer kernel patch.

Solaris x86

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	N/A	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i ¹ - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 X86 kernel patch 144501, or a newer kernel patch.

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
HP-UX 11.31	✓	✓	✓	✓	✗	✗	N/A

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
AIX 6.1	✓	✓	✓	✗	✗	N/A	N/A
AIX 7.1	N/A	✓	✓	✓	✓	✓	N/A
AIX 7.2	N/A	✓	✓	✓	✓	✓	N/A

SQL Server Matrix

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Delphix only supports 64-bit operating systems.

Key:

Color	Supported?
	Yes
	Not Applicable

Windows Server

Supported OS Version	Supported DBMS Version		
	SQL Server 2012	SQL Server 2014	SQL Server 2016
Windows Server 2012			
Windows Server 2012 R2			
Windows Server 2016			
Windows Server 2019			
Windows Server 2022			

SAP ASE Matrix

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• 64-bit OS support only

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
RHEL 5.5				
RHEL 5.6				
RHEL 5.7				
RHEL 5.8				
RHEL 5.9				
RHEL 5.10				

RHEL 5.11	✓	✓	✓	N/A
RHEL 6.0	✗	✗	✗	N/A
RHEL 6.1	✗	✗	✗	✗
RHEL 6.2	✓	✓	✓	✗
RHEL 6.3	✓	✓	✓	✗
RHEL 6.4	✓	✓	✓	✗
RHEL 6.5	✓	✓	✓	✓
RHEL 6.6	✓	✓	✓	✓
RHEL 6.7	✓	✓	✓	✓
RHEL 6.8	✓	✓	✓	✓
RHEL 6.9	✓	✓	✓	✓
RHEL 6.10	✓	✓	✓	✓
RHEL 7.0	N/A	N/A	✓	✓
RHEL 7.1	N/A	N/A	✓	✓
RHEL 7.2	N/A	N/A	✓	✓
RHEL 7.3	N/A	N/A	✓	✓
RHEL 7.4	N/A	N/A	✓	✓
RHEL 7.5	N/A	N/A	✓	✓

RHEL 7.6	N/A	N/A	✓	✓
RHEL 7.7	N/A	N/A	✓	✓
RHEL 7.8	N/A	N/A	✓ 6.0.3+	✓ 6.0.3+
RHEL 7.9	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+
RHEL 8.0	N/A	N/A	N/A	✓
RHEL 8.1	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.2	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.3	N/A	N/A	N/A	✓ 6.0.7+

SUSE Linux Enterprise Server (SLES)

 Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
SLES 11	✓	✗	✗	✗
SLES 11 SP1	✓	✗	✗	✗
SLES 11 SP2	✓	✗	✗	✗
SLES 11 SP3	✓	✗	✗	✗

SLES 11 SP4	✓	✗	✗	✗
SLES 12	✗	✗	✗	✗
SLES 12 SP1	✗	✗	✗	✗
SLES 12 SP2	✗	✗	✗	✗
SLES 12 SP3	✗	✗	✗	✗
SLES 12 SP4	✗	✗	✗	✓ 6.0.4+

Solaris SPARC

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
Solaris 10 U9	✓	✓	✓	✗
Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	✓	✓	✓	✗
Solaris 11	✗	✗	✗	✗
Solaris 11 U1	✗	✗	✗	✗
Solaris 11 U2	✗	✗	✗	✗
Solaris 11 U3	✗	✗	✓	✗
Solaris 11 U4	✗	✗	✓ 6.0.1.0+	✓ 6.0.1.0+

Solaris x86

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16

Solaris 10 U9	✓	✓	✓	✗
Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	N/A	N/A	✓	✗
Solaris 11	N/A	N/A	✓	✓
Solaris 11 U1	N/A	N/A	✓	✓
Solaris 11 U2	N/A	N/A	✓	✓
Solaris 11 U3	N/A	N/A	✓	✓
Solaris 11 U4	N/A	N/A	✗	✗

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
HP-UX 11.31	✗	✗	✗	✗

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
AIX 6.1	✓	✗	✓	✗
AIX 7.1	✗	✗	✓	✓ 6.0.1.0+
AIX 7.2	✗	✗	✓	✓

Db2 Matrix

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			
RHEL 7.1			
RHEL 7.2			

RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3

6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

PostgreSQL Matrix

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	■
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	■
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where x represents the minor version.

 • All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.

• Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

 Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

HANA Matrix

This section describes the HANA (DBMS) versions that are supported by the HANA 2.0 Plugin, as well as the compatible operating systems (OS), for use on the source and target environments.

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Supported DBMS Versions

HANA Version (Standard and Enterprise Editions)	Support Package Stack (SPS)
2.0	04
2.0	05

 **Important:** A target may have a higher SPS release than a source; the reverse is not supported.

Key:

Color	Supported?
	Yes
	No

Supported Operating System

Supported OS Version	Supported DBMS Version	
	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
RHEL 8.1		
RHEL 7.3/7.4/7.5		
CentOS 7.3/7.4/7.5		
RHEL 7.6		

CentOS 7.6	✓	✓
SUSE 12 SP 03	✓	✗
SUSE 15 SP 02	✗	✓

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	Lua-based plugin						vSDK-based Plugin						
	4.1.1	4.1.2	4.2.0	4.3.0	4.3.1	4.3.2	6.0.0	6.0.1	6.1.0	6.1.1	6.2.0	6.3.0	
6.0.0.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.1.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.2.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.3.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.4.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.5.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.6.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.7.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	

	Lua-based plugin						vSDK-based Plugin					
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point-in-Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Although the Staging Push solution is generic and should work with HANA Streaming Replication (HSR) but for this release, we have not validated the same.
- Third-party backups tools are not supported with Pull based ingestion mechanism.
- More than one instance of the same HANA service is not supported with this release.
- Switching between the Staging Push to Pull model is not supported with this release.

EBS Matrix

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		

	Supported EBS 12.2 Apps version	
Supported OS Version	Oracle 11g/12c	Oracle 19c
RHEL 6.7	✓	N/A
RHEL 6.8	✓	N/A
RHEL 6.9	✓	N/A
RHEL 7.0	✓	N/A
RHEL 7.1	✓	N/A
RHEL 7.2	✓	N/A
RHEL 7.3	✓	N/A
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓

	Supported EBS 12.2 Apps version	
Supported OS Version	Oracle 11g/12c	Oracle 19c
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

Reference

Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
1.8.3	
1.8.4	
2.0.0	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.
2.0.1	
3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Related Topics

[EBS 12.2 Plugin Release Notes](#)

Kerberos Support Matrix

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Oracle - Red Hat Enterprise Linux (RHEL)

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Supported OS Version	Supported DBMS Version				
	Oracle 10g	Oracle 11gR1	Oracle 11gR2	Oracle 12cR1	Oracle 12cR2
RHEL 5.5					
RHEL 5.6					
RHEL 5.7					
RHEL 5.8					
RHEL 5.9					
RHEL 5.10					
RHEL 5.11					
RHEL 6.0					

	Supported DBMS Version				
Supported OS Version	Oracle 10g	Oracle 11gR1	Oracle 11gR2	Oracle 12cR1	Oracle 12cR2
RHEL 6.1	N/A	N/A	✓	✓	N/A
RHEL 6.2	N/A	N/A	✓	✓	N/A
RHEL 6.3	N/A	N/A	✓	✓	N/A
RHEL 6.4	N/A	N/A	✓	✓	✓
RHEL 6.5	N/A	N/A	✓	✓	✓
RHEL 6.6	N/A	N/A	✓	✓	✓
RHEL 6.7	N/A	N/A	✓	✓	✓
RHEL 6.8	N/A	N/A	✓	✓	✓
RHEL 6.9	N/A	N/A	✓	✓	✓
RHEL 6.10	N/A	N/A	✓	✓	✓
RHEL 7.0	N/A	N/A	✓	✓	✓
RHEL 7.1	N/A	N/A	✓	✓	✓
RHEL 7.2	N/A	N/A	✓	✓	✓
RHEL 7.3	N/A	N/A	✓	✓	✓
RHEL 7.4	N/A	N/A	✓	✓	✓

	Supported DBMS Version				
Supported OS Version	Oracle 10g	Oracle 11gR1	Oracle 11gR2	Oracle 12cR1	Oracle 12cR2
RHEL 7.5	N/A	N/A	✓	✓	✓

SAP ASE - Red Hat Enterprise Linux (RHEL)

	Supported DBMS Version			
Supported OS Version	ASE 15.0.3	ASE 15.5	ASE 15.7	ASE 16
RHEL 5.5	✓	✓	✓	N/A
RHEL 5.6	✓	✓	✓	N/A
RHEL 5.7	✓	✓	✓	N/A
RHEL 5.8	✓	✓	✓	N/A
RHEL 5.9	✓	✓	✓	N/A
RHEL 5.10	✓	✓	✓	N/A
RHEL 5.11	✓	✓	✓	N/A
RHEL 6.0	✗	✗	✗	N/A
RHEL 6.1	✗	✗	✗	✗
RHEL 6.2	✓	✓	✓	✗
RHEL 6.3	✓	✓	✓	✗
RHEL 6.4	✓	✓	✓	✗

Supported OS Version	Supported DBMS Version			
	ASE 15.0.3	ASE 15.5	ASE 15.7	ASE 16
RHEL 6.5	✓	✓	✓	✓
RHEL 6.6	✓	✓	✓	✓
RHEL 6.7	✓	✓	✓	✓
RHEL 6.8	✓	✓	✓	✓
RHEL 6.9	✓	✓	✓	✓
RHEL 6.10	✓	✓	✓	✓
RHEL 7.0	N/A	N/A	✓	✓
RHEL 7.1	N/A	N/A	✓	✓
RHEL 7.2	N/A	N/A	✓	✓
RHEL 7.3	N/A	N/A	✓	✓
RHEL 7.4	N/A	N/A	✓	✓
RHEL 7.5	N/A	N/A	✓	✓

DB2 Support Matrix



- ESE: Enterprise Server Edition
- AESE: Advanced Enterprise Server Edition
- For the supported DB2 versions, Delphix supports the corresponding DB2 Developer edition where the vendor, IBM, supports it
- 64-bit OS support only

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported DBMS Version			
	ESE 10.5	AESE 10.5	ESE 11.1	AESE 11.1
RHEL 6.0	✗	✗	✗	✗
RHEL 6.1	✗	✗	✗	✗
RHEL 6.2	✗	✗	✗	✗
RHEL 6.3	✗	✗	✗	✗
RHEL 6.4	✓	✓	✗	✗
RHEL 6.5	✓	✓	✗	✗
RHEL 6.6	✓	✓	✗	✗
RHEL 6.7	✓	✓	✓	✓
RHEL 6.8	✓	✓	✓	✓
RHEL 6.9	✓	✓	✓	✓
RHEL 7.0	✓	✓	✓	✓
RHEL 7.1	✓	✓	✓	✓
RHEL 7.2	✓	✓	✓	✓
RHEL 7.3	✓	✓	✓	✓
RHEL 7.4	✓	✓	✓	✓

 DB2 11.1 Developer Edition has been certified on RHEL 6.7.

DB2 on Advanced Interactive eXecutive (AIX)

Supported OS Version	Supported DBMS Version			
	ESE 10.5	AESE 10.5	ESE 11.1	AESE 11.1
AIX 6.1	✘	✘	N/A	N/A
AIX 7.1	✔	✔	✔	✔
AIX 7.2	✔	✔	✔	✔

Data Source Certifications

- [Legal Notice](#)
- [Goals](#)
 - [Certification Expectations:](#)
- [Data Sources](#)
- [Cadence for Data Source Database Version Changes](#)
- [Cadence for Operating System Version Changes](#)
- [Appendix](#)
 - [Product Lifecycle of Oracle Related DB/OS Vendors:](#)
 - [Product Lifecycle of SQL Server related DB/OS vendors:](#)
 - [Product Lifecycle of SAP ASE related DB/OS vendors:](#)
 - [Footnote 1: Update Version Terminology](#)

Legal Notice

IMPORTANT: The policies outlined in this document are subject to change. The policies and timelines are provided as rough estimates when certifications of interoperability/connectivity with third-party products may be completed and are subject to whether such interoperability/connectivity is part of the Delphix product roadmap and support policy and should not be regarded as a commitment.

Speak with Delphix Product Management if you need more specific information or are unsure of these policies.

Goals

The goal of this certification cadence is to maintain a cadence for certifications of new software releases.

The objectives are:

1. To set reasonable expectations for our organization and customers regarding how and when we will address certifications
2. To define an intended timeframe for certification of new data sources versions and operating systems versions to maintain an up-to-date support matrix

Certification Expectations:

Please note that certification only covers previously supported functionality of the Delphix Engine. Any new features and functionality introduced in the database or operating system (DB/OS) will be tracked through the regular release planning and roadmap process and considered as features/functionality/fixes.

Specifically, the following paths exist with varying delivery timeline commitments, in the following priority order:

1. [Security Response Policy](#): In case security issues are discovered in a certain DB and/or OS version, we will follow our security response policy, which supersedes the goals described in this document.
2. [Certifications](#): As described in the document.
3. [Roadmap Features](#): No policies outlined. This follows the regular roadmap planning process.

Data Sources

This documentation currently applies only to the following data source integrations:

- Oracle
- SQL Server
- SAP ASE

Cadence for Data Source Database Version Changes

Release Type	Certification Details	Goal
A new version of a database with all currently supported operating systems	Certification with all supported operating systems	8 months from new version GA date

Cadence for Operating System Version Changes

Release Type	Certification Details	Goal
A new major operating system version with all currently supported databases	Certification with all supported versions of the database	6 months from new version GA date
A new update* operating system version with all currently supported databases	Certification with all supported versions of the database	4 months from new version GA date

*Terminology for an 'update' version varies across operating systems. See the Appendix for more information.

Appendix

Product Lifecycle of Oracle Related DB/OS Vendors:

Software	Product Lifecycle	Notes
Oracle Database	Link	12.1 is Oracle Enterprise Edition (EE)
Red Hat Enterprise Linux	Link	
Solaris (x86 and SPARC)	Link	
SuSE Linux (SLES)	Link	
IBM AIX	Link	
HP-UX	Link	

Product Lifecycle of SQL Server related DB/OS vendors:

Software	Product Lifecycle
Microsoft SQL Server	Link

Microsoft Windows	Link
-------------------	----------------------

Product Lifecycle of SAP ASE related DB/OS vendors:

Software	Product Lifecycle
SAP ASE	Link
Red Hat Enterprise Linux	Link
Solaris (x86 and SPARC)	Link
SuSE Linux (SLES)	Link
IBM AIX	Link

Footnote 1: Update Version Terminology

Operating System releases typically include updated versions, the terminology for which varies depending on the vendor. Below are the various operating systems which we support, as well as the terminology for updates to use with the document above.

Vendor	Update Version Terminology	Definition
Red Hat Enterprise Linux	Update version	Installed on top of major releases to provide larger-scale point-release updates via Red Hat Network, after initial installation via traditional methods (PXE, DVD, etc.).
Solaris (x86 and SPARC)	Update release	Update release number. The update release number for this Oracle Solaris release. The update value is 0 for the first customer shipment of an Oracle Solaris release, 1 for the first update of that release, 2 for the second update of that release, and so forth.
SuSE Linux (SLES)	Service Pack (SP) release	Service Packs are minor updates and enhancements, usually with security fixes, that are applied to the major version of the OS.

IBM AIX	Service Pack (SP) release	A Technology Level (TL) gives new features, while an SP contains fixes for problems that are critical and can't wait until the next TL. Service Packs are limited to minimal corrections that don't change the way things work or add any new functionality.
HP-UX	Patch release	Patch sets are groups of patches that should be installed to support common HPE products and sub-systems.
Windows	New Release (R) version or Service Pack (SP) release	A service pack (SP) is a collection of updates and fixes, called patches, for an operating system or a software program. Many of these patches are often released before a larger service pack, but the service pack allows for an easy, single installation.

Virtualization Platform Matrix

Legend

Color	Supported?
✓	Yes
✗	No
N/A	Not Applicable

Delphix Features with Data Platforms

This section describes the Delphix features we support for each data platform.

Delphix Feature		Oracle						SQL Server					SAP ASE			
		11g	12c	18c	19c	PD B	21c	2008	2012	2016	2017	2019	15.3	15.5	15.7	16
Environment Management	Manual environment discovery	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Automated environment discovery	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Hostchecker	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Environment monitor	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Linking/ Syncing	Load from existing backup	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Incremental snapshots	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Resumable syncs for initial load	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
	dSource LogSync	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Linking clustered source databases	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
	Validated Sync on snapshot	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Validated Sync, continuous	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Load from alternate sources	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
	VDB LogSync	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗

	Source Continuity	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
	Handling of source DB upgrades	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
	Handling of non-logged changes	✓	✓	✓	✓	✓	✓	N/A	N/A	N/A	N/A	N/A	✗	✗	✗	✗
Provisioning	VDB on VDB	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	VDB refresh and rewind	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Provisioning clustered VDBs	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
	VDB migration	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	VDB upgrades outside of Delphix	✓	✓	✓	✓	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓

	Automated VDB upgrades with Delphix	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Database config parameters for VDBs	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Path-mapping	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗
V2P for Multi tenant	Basic V2P data export	N/A	N/A	N/A	N/A	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	V2P over DSP	N/A	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Multi filesystem V2P	N/A	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Resumable V2P	N/A	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
V2P for Non-Multi tenant	Basic V2P data export	✓	✓	✓	✓	N/A	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	V2P over DSP	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
	Multi filesystem V2P	✓	✓	✓	✓	N/A	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Resumable V2P	✓	✓	✓	✓	N/A	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

Data Platform Features

This section describes the data platform-specific features we support for each use case within Delphix.

*Non-Multitenant TDE support does not apply to vPDBs.

 Multitenant TDE support applies to vPDBs starting with Oracle version 12cR2 (12.2) and above.

*Oracle databases can have some tables that are HCC-enabled and some that are not; Delphix can ingest a database with HCC-enabled tables and can successfully provision a VDB. The caveat to this is that attempts to read data from the HCC-enabled tables on a provisioned VDB will fail. This happens because Oracle requires the backing storage to be fully supported for HCC, but uncompressed tables can be used as expected with non-HCC tables. Please note that Oracle only recommends HCC for tables or partitions with little to no DML operations. Another possible solution is creating hook operations to uncompress the HCC-enabled tables/partitions to access the data, this method is outlined in a [Delphix Community](#) post.

Oracle

Platform Feature		Oracle					
		11g	12cR1 (12.1)	12cR2 (12.2)	18c	19c	21c
dSource Linking	Non-Multitenant	✓	✓	✓	✓	✓	N/A
	CDB with PDB Multitenant	N/A	✓	✓	✓	✓	✓
	Read-Only Oracle Home	✗	✗	✗	✓	✓	✓
	CDB with Application Container PDB	✗	✗	✗	✗	✗	✗
dSource SnapSync	Level based backups	✓	✓	✓	✓	✓	✓
	SCN based backups	✓	✓	✓	✓	✓	✓
	Physical Standby	✓	✓	✓	✓	✓	✓
	Active Dataguard	✓	✓	✓	✓	✓	✓
	ASM	✓	✓	✓	✓	✓	✓

	RAC	✓	✓	✓	✓	✓	✓
	TDE (non-multitenant)*	✓	✓	✓	✓	✓	✓
	TDE (multitenant)	N/A	✗	✓	✓	✓	✓
	HCC*	✗	✗	✗	✗	✗	✗
	Exadata	✓	✓	✓	✓	✓	✓
dSource LogSync	RAC	✓	✓	✓	✓	✓	✓
	ASM FS	✓	✓	✓	✓	✓	✓
	Native FS	✓	✓	✓	✓	✓	✓
	Physical Standby	✓	✓	✓	✓	✓	✓
VDB	Non-Multitenant	✓	✓	✓	✓	✓	N/A
	RAC	✓	✓	✓	✓	✓	✓
	Read-Only Oracle Home	✗	✗	✗	✓	✓	✓
	Application Container	✗	✗	✗	✗	✗	✗

SQL Server

Platform Feature		SQL Server				
		2008	2012	2016	2017	2019
Linking and Syncing	Linking from clustered sources	✓	✓	✓	✓	✓
	Full and bulk-logged mode support	✓	✓	✓	✓	✓
	Simple mode support	✓	✓	✓	✓	✓

	Filestreams	✓	✓	✓	✓	✓
	Full textcatalogs	✓	✓	✓	✓	✓
	Availability groups	N/A	✓	✓	✓	✓
Backup Software	Native backups	✓	✓	✓	✓	✓
	Litespeed SQL Server backups	✓	✓	✓	✓	✓
	Red Gate SQL Backup Pro	✓	✓	✓	✓	✓
	NetBackup	✓	✓	✓	✓	✓
	Commvault	✓	✓	✓	✓	✓
Provisioning	Clustered VDBs	✓	✓	✓	✓	✓
	Availability Groups	N/A	✗	✗	✗	✗
	Change Data Capture	✓	✓	✓	✓	✓

SAP ASE

Platform Feature		SAP ASE			
		15.0.3	15.5	15.7	16
Linking and Syncing	Linking from clustered sources	✗	✗	✗	✗
	Linking from ASE replication server	✗	✗	✗	✗
	Full mode support	✓	✓	✓	✓
	Truncate log on checkpoint support	✓	✓	✓	✓

	Native backups		✓	✓	✓	✓
	NetBackup		✗	✗	✗	✗
Provisioning	Clustered VDBs		✗	✗	✗	✗
	VDB LogSync		✗	✗	✗	✗

vFiles Matrix

Unix Environments

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Color	Supported?
	Yes

Red Hat Enterprise Linux (RHEL)

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
RHEL 5.5	
RHEL 5.6	
RHEL 5.7	
RHEL 5.8	
RHEL 5.9	
RHEL 5.10	
RHEL 5.11	
RHEL 6.0	
RHEL 6.1	
RHEL 6.2	
RHEL 6.3	

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
RHEL 6.4	✓
RHEL 6.5	✓
RHEL 6.6	✓
RHEL 6.7	✓
RHEL 6.8	✓
RHEL 6.9	✓
RHEL 6.10	✓
RHEL 7.0	✓
RHEL 7.1	✓
RHEL 7.2	✓
RHEL 7.3	✓
RHEL 7.4	✓
RHEL 7.5	✓
RHEL 7.6	✓
RHEL 7.7	✓
RHEL 7.8	✓
RHEL 7.9	✓ 6.0.7+

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
RHEL 8.0	✓
RHEL 8.1	✓ 6.0.3+
RHEL 8.2	✓ 6.0.3+
RHEL 8.3	✓ 6.0.8+
RHEL 8.4	✓ 6.0.11.0+

Oracle Enterprise Linux

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
Oracle Enterprise Linux 5.5	✓
Oracle Enterprise Linux 5.6	✓
Oracle Enterprise Linux 5.7	✓
Oracle Enterprise Linux 5.8	✓
Oracle Enterprise Linux 5.9	✓
Oracle Enterprise Linux 5.10	✓
Oracle Enterprise Linux 5.11	✓
Oracle Enterprise Linux 6.0	✓
Oracle Enterprise Linux 6.1	✓

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
Oracle Enterprise Linux 6.2	✓
Oracle Enterprise Linux 6.3	✓
Oracle Enterprise Linux 6.4	✓
Oracle Enterprise Linux 6.5	✓
Oracle Enterprise Linux 6.6	✓
Oracle Enterprise Linux 6.7	✓
Oracle Enterprise Linux 6.8	✓
Oracle Enterprise Linux 6.9	✓
Oracle Enterprise Linux 6.10	✓
Oracle Enterprise Linux 7.0	✓
Oracle Enterprise Linux 7.1	✓
Oracle Enterprise Linux 7.2	✓
Oracle Enterprise Linux 7.3	✓
Oracle Enterprise Linux 7.4	✓
Oracle Enterprise Linux 7.5	✓
Oracle Enterprise Linux 7.6	✓
Oracle Enterprise Linux 7.7	✓

	Delphix AppData Version (x86_64)
Supported OS Version	6.0+
Oracle Enterprise Linux 7.8	✓
Oracle Enterprise Linux 7.9	✓ 6.0.7+
Oracle Enterprise Linux 8.0	✓
Oracle Enterprise Linux 8.1	✓
Oracle Enterprise Linux 8.2	✓ 6.0.5+

SUSE Linux Enterprise Server (SLES)

Delphix AppData Version	6.0+
Supported OS Version (x86_64)	
SLES 11	✓
SLES 11 SP1	✓

Solaris

Delphix AppData Version	6.0.+
Supported OS Version (SPARC x86_64)	
Solaris 10	✓
Solaris 11	✓

AIX

Solaris

Delphix AppData Version	6.0.+
--------------------------------	--------------

	Supported OS Version (Power)	
	AIX 6.1	✓
	AIX 7.1	✓
	AIX 7.2	✓

HP-UX

Solaris

Delphix AppData Version	6.0+
Supported OS Version (IA64)	
HP-UX 11i v3 (11.31)	✓

Windows Environments

Delphix AppData Version	6.0+
Supported OS Version	
Windows Server 2012	✓
Windows Server 2012 R2	✓
Windows Server 2016	✓
Windows Server 2019	✓
Windows Server 2022	✓

Windows Connector Matrix

You can always download the current Windows connector version from the Delphix Engine from the following locations:

Delphix Dynamic Data Platform versions before 6.0.5.0:

```
http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.msi
```

Delphix Dynamic Data Platform versions 6.0.5.0 onwards:

```
http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.exe
```

Note

The Delphix Connector and Engine are generally backward-compatible. Delphix connector version 1.18.0 and beyond is compatible with Delphix Engine 6.0.5.0 and later releases. Delphix connector version 1.17.0 or earlier is compatible with Delphix Engine 6.0.4.0 and earlier releases.

Checking the Windows Connector Version

After installing the connector, you can perform the following steps to determine the Windows connector version.

1. Navigate to the Windows Environment.
2. In the search box on the taskbar, type Add or Remove Programs or launch **Start > Control Panel > Programs > Programs and Features**.
3. On the Add or Remove Programs screen, locate Delphix Connector and then click on the software to check the version installed.

Windows Connector Matrix

The following table details the supported Windows Connector and JRE versions for the corresponding Delphix Dynamic Data Platform versions.

Delphix Dynamic Data Platform Version	Windows Connector Version	JRE Version
4.2	1.2.0	1.7.0_71 (4.2.5.1)
4.3	1.3.0	1.8.0_40 (4.3.5.1)
5.1.0.0-5.1.4.0	1.4.0	1.8.0_40
5.1.5.0-5.1.9.0	1.5.0	1.8.0_112
5.2.0.0-5.2.2.0	1.9.0	1.8.0_131 (5.2.2.1)

Delphix Dynamic Data Platform Version	Windows Connector Version	JRE Version
5.2.3.0-5.2.4.0	1.10.0	1.8.0_162
5.2.5.0-5.3.3.0	1.12.0	1.8.0_202
5.3.4.0	1.13.0	1.8.0_202
5.3.5.0	1.14.0	1.8.0_202
5.3.6.0-5.3.7.0	1.15.0	1.8.0_202
5.3.8.0-6.0.0.0	1.16.0	1.8.0_202
6.0.1.0-6.0.4.0	1.17.0	OpenJDK 8u242
6.0.5.0-6.0.8.0	1.18.0	OpenJDK 8u262-b10
6.0.9.0-6.0.10.0	1.20.0	OpenJDK 8u282-b08
6.0.11.0	1.21.0.0	OpenJDK 8u302-b08
6.0.12.0	1.22.0.0	OpenJDK 8u302-b08
6.0.12.1	1.22.0.0	OpenJDK 8u302-b08
6.0.13.0	1.23.0.0	OpenJDK 8u302-b08

The following table details the .Net connector version compatibility for the corresponding Windows Connector versions.

.Net Connector Version	Windows Connector Version
3.5	1.17.0 and earlier
3.5 and 4.0+	1.18.0 and higher

File Checksum / Hashes

The following table provides the MD5 and SHA256 file hash for recent Delphix versions, in the instance that a security policy requires an exception for software installation.

Delphix Dynamic Data Platform Version	Connector Version	MD5 Hash	SHA256 Hash
6.0.0.0	1.16.0	84616c4f86c0d871127feb26b9b8bd52	bc2e5e4b26c33d739f7be12f53637979481e1371325e0b5c7645462f6cf5f4f3
6.0.1.0	1.17.0	bb6544dcff3f22d6bf9e20429b1cdbc0c	8d914b77d2c5aa0695c3fdb7d7394e8ae051c091f6f5ee59185cb1d9d157542
6.0.1.1	1.17.0	dd3ea3bedf92a1f8d7e89228315ee315	75ed8aa2cf2689916e0b92de6c2aa2fa7660003d62ff07e99978f95dcc9f6abc
6.0.2.0	1.17.0	4770f8c1ad46430a1e4b8764ad687ff5	a98f1c1f90d5053f62d74c1a74bc27b3b48f68a3e067dce04ce9ab2ccf484241
6.0.2.1	1.17.0	97ec3d13db50ad843625b7757c1a2360	7c2ab497ae4ac27b907611e976d7068ed8eaf54e52b836591c082bbe01079512
6.0.3.0	1.17.0	3154c7a5fd7bcb4f0e2c66d328290e6a	efb69b2df884705e488d3405a3958c3bb5457142cbfc71990559bb549c2e8c37
6.0.3.1	1.17.0	2f73be2cf8ceef97cf519c716237462d	7fa7bc269bbc944568231f93ac17ba919b494aa43009967245b75f4abe576a73
6.0.4.0	1.17.0	6a4b7f0a5d180438ff199b2b464506cd	971b88456574e8e11906b5e6850ea46f012490dcb789a5af73f37d71158ffebf
6.0.4.1	1.17.0	ea9a564d0cae1066cd9607824292fbd7	674f9896360a351b314d6aa6e624f9270ef9e3e8d113baddf028581f1ddddec9
6.0.4.2	1.17.0	c5f855c1fa26f5cd3d84d70a94fc0480	07d71acafd559e15782e7145225734adf5e50356d2f8abe5a52bbd90f3a44cbe
6.0.5.0	1.18.0	a3ac768989a6c1dae41adf9452548236	71525c31edcf321d2c0220f4edf8932e249e2cd14b261054af1859574dc9e5eb
6.0.6.0	1.18.0	acb09cbd754e9c25d1afa17b35321cd4	3a51bd3b05148dc3563c5b36d83d44ac725295c23c4b7e9c28f54263058fe2b4
6.0.6.1	1.18.0	32202760adb743a62851f8b698a31bf3	ac16e8a5c37807c2fa775ecc18f9e0387f5bd9d235c9e8854275c99d5bf865a1

Delphix Dynamic Data Platform Version	Connector Version	MD5 Hash	SHA256 Hash
6.0.7.0	1.18.0	2fed8da4b986b95a7efb9a9a148239ba	2c2a83bad568617bf5dd36b84ad0de22b811f113acac45124ed0effbeb6aef17
6.0.8.0	1.18.0	5bf0b950cccad8c39db60e1a7406d985	eae5958878a22b4f4f0456210c031ef9c11d064a05ca6cb85535f0e7dec9c0a1
6.0.8.1	1.18.0	a3ac768989a6c1dae41adf9452548236	71525c31edcf321d2c0220f4edf8932e249e2cd14b261054af1859574dc9e5eb
6.0.9.0	1.20.0	127af1d95372f84feb373ab4f7b280b2	b7c7aaceb2b04aebfbd0ba2521dd15520159d54877b0543567b892ee2584c7ce
6.0.10.0	1.20.0	e8a39a57e5a77d106d8040a4537bd8f0	7d4501bf5f1cf22a7039e8f7429d4b42c0d11b1360434380b174830f1155646e
6.0.10.1	1.20.0	baf1223b73e626798dd404bd4d18a3c9	b827d9213e933fce8389322fe25c8d2e1db57ab66709d236b9a8cd5e254bb89d
6.0.11.0	1.21.0.0	412c4c32c7aefb973c21b573dca66d2a	9b689a21249456ddf623042cd2595820531cfbfca0703e713d9795036eef49b7
6.0.12.0	1.22.0.0	dc7a97c46ef9e2d91285dcb0712eaf61	db544774ac3eb3235bf90deef3e93962f3955e86bd9f14309ecb0fba9d41dd7f
6.0.12.1	1.22.0.0	8cf9e955438c6393031c5596ca0420cd	130249658f68eab2a2263deceed6a9de6dd54eead66da68886272f3f40b2f9af
6.0.13.0	1.23.0.0	31cc5d3869e4c7f93d38014e5bd40b27	de1c2fd4c79dcb7391a51774fba1821456ecc031283db5b497145cce383005f2

Migration and Compatibility

This section covers the following topics:

- [Upgrade Matrix](#)
- [Supported Oracle DBMS Versions and Operating Systems for Source and Target Environments](#)
- [Supported SAP ASE Versions and Operating Systems](#)
- [Supported SQL Server Versions, Operating Systems, and Backup Software](#)
- [Tested Browser and Operating Systems](#)

Upgrade Matrix

6.0.x Upgrade Matrix

This section lists all the Delphix Engine versions that a user can upgrade to the required or the latest version.

i Self-service upgrades to the Delphix Engine 6.0.* from 5.3.* are not supported. 6.0.6.1 is the last version that can be directly upgraded to, from 5.3.6.0-5.3.9.0.

! Starting with 6.0.10.0, the policy of supporting direct upgrades from software released in the last 12 months will be enforced. Only 6.0.4.0 and above can upgrade directly to 6.0.10.0. If 6.0 engines exist that are older than 6.0.4.0, they need to be upgraded to 6.0.7.0 first.

Please refer to the documentation on [How to Upgrade](#) or check with Delphix Support to perform the upgrade.

	Upgrade To																				Upgrade From	
	6.0.0.0	6.0.0.1	6.0.0.2	6.0.0.3	6.0.0.4	6.0.0.5	6.0.0.6	6.0.0.7	6.0.0.8	6.0.0.9	6.0.0.10	6.0.0.11	6.0.0.12	6.0.0.13	6.0.0.14	6.0.0.15	6.0.0.16	6.0.0.17	6.0.0.18	6.0.0.19		
5.3.5.0 and prior	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	5.3.5.0 and prior
5.3.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	x	x	x	5.3.6.0

5.3.7.0	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	5.3.7.0
5.3.7.1	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	5.3.7.1
5.3.8.0	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	5.3.8.0
5.3.8.1	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	5.3.8.1
5.3.9.0	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	5.3.9.0
6.0.0.0	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	6.0.0.0
6.0.1.0	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	6.0.1.0
6.0.1.1	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	6.0.1.1
6.0.2.0	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	6.0.2.0

6.0.2.1	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	x	6.0.2.1			
6.0.3.0	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	x	6.0.3.0			
6.0.3.1	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	x	6.0.3.1			
6.0.4.0	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.4.0		
6.0.4.1	x	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.4.1		
6.0.4.2	x	x	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.4.2		
6.0.5.0	x	x	x	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.5.0	
6.0.6.0	x	x	x	x	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.6.0
6.0.6.1	x	x	x	x	x	x	x	x	x	x	x	x	x	✓	✓	✓	✓	✓	✓	✓	✓	x	x	x	x	x	6.0.6.1

Supported Oracle DBMS Versions and Operating Systems for Source and Target Environments

✔ Source and Target OS and DBMS Compatibility

- Source and Target must have the same major Oracle version and Edition (Enterprise/Standard). It is recommended to also have the same minor versions and patches.
 - If the minor version (specified with the second digit with Oracle version 18+, or the dated patch level with Oracle version 12.2) or patches differ between the source and target, then a plug-in violation will be received when opening the vPDB on the target, and will need to run datapatch manually or via a hook to resolve it.
- Source and Target operating systems must be compatible. E.g., AIX -> AIX, Solaris SPARC -> Solaris SPARC, x86_64 Linux -> x86_64 Linux.
- In addition, x86_64 Solaris -> x86_64 Linux and x86_64 Linux -> x86_64 Solaris is supported.

⚠ Oracle 10.2.0.4

The Delphix Engine does not support Oracle 10.2.0.4 databases using Automatic Storage Management (ASM) that do not have the patch set for Oracle Bug 7207932. This bug is fixed in patch set 10.2.0.4.2 onward.

Disclaimers and Support Notes

Delphix Support Policies specifically list **major** and **minor** release coverage. If a minor release is listed as covered, all patch releases under that minor release are certified. Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL), 64-bit OS support only.

The Oracle version in the support matrix applies to both Oracle DB and Oracle Grid. Delphix supports the same set of features and functionality for supported non-multitenant database versions. Currently, Delphix does not support the Oracle 12c feature of THREADED_EXECUTION being set to TRUE, because this disables OS authentication.

Summary of Delphix features that are unsupported for the Oracle12c and higher Multi-tenant configuration:

- Existing Virtual Container Database (vCDB) as target for provisioning an additional vPDB
- Customize VDB settings/initialization parameters. Includes the following:
 - Customize init.ora database parameter during provisioning
 - Config templates
 - Online redo log size
 - Number of RAC VDB instances
 - Online redo log groups
 - Archive log mode
 - Setting new DBID
 - Customize local listeners
- Virtual to Physical (V2P) Support
- Resumable initial SnapSync
- Validated Sync
- Source continuity for dSource upgraded from Oracle 12c non-multitenant to multitenant database
- Cross-platform provisioning (XPP) to the virtual database
- Oracle LiveSources

Delphix supports systems that are part of an Exadata or Exadata Cloud-at-Customer (ExaCC) cluster, provided that the operating system and DBMS version are in the supported list. Oracle Exadata, Exadata Cloud, and Exadata Cloud-at-Customer (ExaCC) offer several features that are reliant on the underlying Exadata storage. Delphix

supports creating dSources from Exadata source clusters and provisioning VDBs and vPDBs to Exadata target clusters.

The following Exadata-only features will not be available in child VDBs running in Delphix storage:

- Smart Scan - allows SQL processing to be offloaded to Exadata storage cells. Queries that leverage Smart Scan will execute successfully in VDBs but query performance may be impacted.
- Smart Flash Cache - allows frequently used data blocks to be stored in the flash cache of Exadata storage cells to allow for faster access. Queries against objects that are normally cached in Smart Flash Cache will execute successfully in VDBs but query performance may be impacted. As the virtualization engine uses memory to cache frequently used data blocks, allocating extra memory to the engine may help to mitigate this performance impact.
- Hybrid Columnar Compression (HCC) - compresses data and stores it in a columnar format, allowing Exadata storage cells to deliver faster query performance while also saving on storage. Columns compressed using HCC cannot be used by queries executed in VDBs. Affected columns can be decompressed though care should be taken to evaluate the impact of doing so on VDB refresh times and Delphix storage usage. In Delphix VDBs, queries against partitions or tables using HCC will fail with the following error "ORA-64307: Exadata Hybrid Columnar Compression is not supported for tablespaces on this storage type".

Color	Supported?
	Yes
	No
	Not Applicable

Oracle Database Editions

Delphix supports the following Oracle database editions:

- Enterprise
- Standard

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please refer to [KBA 5622](#) for the required steps.

Supported OS Version	Supported DBMS Version						
	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 5.5							
RHEL 5.6							

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 5.7	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.8	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.9	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.10	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 5.11	✓	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.0	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.1	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.2	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.3	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.4	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.5	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.6	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.7	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.8	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.9	N/A	✓	✓	✓	✓	N/A	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 6.10	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.0	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.1	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.2	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.3	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.4	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.5	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.6	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.7	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.8	N/A	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	N/A
RHEL 7.9	N/A	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.5+	✓ 6.0.4+	N/A
RHEL 8.0	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.1	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.2	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 8.3	N/A	N/A	N/A	N/A	N/A	✓ 6.0.7+	✓ 6.0.11+
RHEL 8.4	N/A	N/A	N/A	N/A	N/A	✓ 6.0.10+	✓ 6.0.11+

SUSE Linux Enterprise Server (SLES)

 Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 11	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP1	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP2	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP3	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP4	✗	✓	✓	N/A	N/A	N/A	N/A
SLES 12	N/A	✓	✓	N/A	N/A	N/A	N/A
SLES 12 SP1	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP2	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP3	N/A	✓	✓	✓	✓	✓	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 12 SP4	N/A	✓	✓	✓	✓	✓	N/A
SLES 12 SP5	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+	✗	✓ 6.0.4+	N/A
SLES 15	N/A	✗	✗	✓	✓	✓	N/A
SLES 15 SP1	N/A	✗	✗	✗	✗	✓	✓ 6.0.11+
SLES 15 SP2	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11+
SLES 15 SP3	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11+

Solaris SPARC

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	✗	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i ¹ - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 SPARC kernel patch 144500, or a newer kernel patch.

Solaris x86

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	N/A	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i **1** - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 X86 kernel patch 144501, or a newer kernel patch.

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
HP-UX 11.31	✓	✓	✓	✓	✗	✗	N/A

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
AIX 6.1	✓	✓	✓	✗	✗	N/A	N/A
AIX 7.1	N/A	✓	✓	✓	✓	✓	N/A
AIX 7.2	N/A	✓	✓	✓	✓	✓	N/A

Delphix supports all x86/x64 CPU environments for source and target.

⚠ Required HP-UX patch for Target Servers

PHNE_37851 - resolves a known bug in HP-UX NFS client prior to HP-UX 11.31.

Supported SAP ASE Versions and Operating Systems

This topic describes supported operating systems and database versions for SAP ASE.

Supported DBMS Versions

- ASE 15.0.3
- ASE 15.5
- ASE 15.7
- ASE 16.0

Supported Operating Systems

✔ Source and Target OS and DBMS Compatibility

The source and target must be running the same DBMS/Operating System combination, (*although users can run different patch/sp levels*) in order to successfully provision a VDB to the target. For example, if the source is running SAP ASE 16, the target can be running ASE 16SP1. However, for ASE 15.7 the versions of the source and target should both be in the 15.7 SP100 to 15.7 SP140 or 15.7 ESD#1 to 15.7 SP64 ranges. There can be no mixing of versions between these two ranges. If the target is used as a staging server, the same rule applies (prior to Delphix 5.2.5.0, the staging server had to match the source server's version down to the SP level). The Operating System platform must be the same between the source and target, even when the operating system version may differ. For example, if the source is running Red Hat Enterprise Linux 6.2 x86_64 then the target could be running Red Hat Enterprise Linux 6.4 x86_64, but not Solaris 10 SPARC.

ASE itself may or may not allow the **MOUNT** command to work between versions. Be sure to check that the **UNMOUNT** and **MOUNT** commands work between your desired ASE versions.

So for example, if we upgraded the ASE staging instance to ASE 15.7 SP140 but we keep the ASE instance hosting the VDBs at ASE 15.7 SP135, we would want to make sure that we can **UNMOUNT** a database from ASE 15.7 SP140 and then **MOUNT** it on ASE 15.7 SP135:

```
isql -Usa -Psybase -SASE157SP140
1> UNMOUNT DATABASE testdb to "/tmp/manifest"
2> GO
```

```
isql -Usa -Psybase -S ASE157SP135
1> MOUNT DATABASE testdb from "/tmp/manifest"
2> go
Msg 14580, Level 16, State 1:
Server 'ASE157SP135', Line 1:
You cannot mount the database(s) because at least one database contains
functionality that is available only on the server on which it originated.
```

In the above example, ASE would not allow the database "testdb" to be mounted on the older release of ASE.

Delphix supports all 64-bit OS environments for source and target.

⚠ Required HP-UX patch for Target Servers

PHNE_37851 - resolves a known bug in HP-UX NFS client prior to HP-UX 11.31.

Support Matrix



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- 64-bit OS support only

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)



Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
RHEL 5.5				
RHEL 5.6				
RHEL 5.7				
RHEL 5.8				
RHEL 5.9				
RHEL 5.10				
RHEL 5.11				

RHEL 6.0	✗	✗	✗	N/A
RHEL 6.1	✗	✗	✗	✗
RHEL 6.2	✓	✓	✓	✗
RHEL 6.3	✓	✓	✓	✗
RHEL 6.4	✓	✓	✓	✗
RHEL 6.5	✓	✓	✓	✓
RHEL 6.6	✓	✓	✓	✓
RHEL 6.7	✓	✓	✓	✓
RHEL 6.8	✓	✓	✓	✓
RHEL 6.9	✓	✓	✓	✓
RHEL 6.10	✓	✓	✓	✓
RHEL 7.0	N/A	N/A	✓	✓
RHEL 7.1	N/A	N/A	✓	✓
RHEL 7.2	N/A	N/A	✓	✓
RHEL 7.3	N/A	N/A	✓	✓
RHEL 7.4	N/A	N/A	✓	✓
RHEL 7.5	N/A	N/A	✓	✓
RHEL 7.6	N/A	N/A	✓	✓

RHEL 7.7	N/A	N/A	✓	✓
RHEL 7.8	N/A	N/A	✓ 6.0.3+	✓ 6.0.3+
RHEL 7.9	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+
RHEL 8.0	N/A	N/A	N/A	✓
RHEL 8.1	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.2	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.3	N/A	N/A	N/A	✓ 6.0.7+

SUSE Linux Enterprise Server (SLES)

 Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
SLES 11	✓	✗	✗	✗
SLES 11 SP1	✓	✗	✗	✗
SLES 11 SP2	✓	✗	✗	✗
SLES 11 SP3	✓	✗	✗	✗
SLES 11 SP4	✓	✗	✗	✗

SLES 12	✗	✗	✗	✗
SLES 12 SP1	✗	✗	✗	✗
SLES 12 SP2	✗	✗	✗	✗
SLES 12 SP3	✗	✗	✗	✗
SLES 12 SP4	✗	✗	✗	✓ 6.0.4+

Solaris SPARC

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
Solaris 10 U9	✓	✓	✓	✗
Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	✓	✓	✓	✗
Solaris 11	✗	✗	✗	✗
Solaris 11 U1	✗	✗	✗	✗
Solaris 11 U2	✗	✗	✗	✗
Solaris 11 U3	✗	✗	✓	✗
Solaris 11 U4	✗	✗	✓ 6.0.1.0+	✓ 6.0.1.0+

Solaris x86

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
Solaris 10 U9	✓	✓	✓	✗

Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	N/A	N/A	✓	✗
Solaris 11	N/A	N/A	✓	✓
Solaris 11 U1	N/A	N/A	✓	✓
Solaris 11 U2	N/A	N/A	✓	✓
Solaris 11 U3	N/A	N/A	✓	✓
Solaris 11 U4	N/A	N/A	✗	✗

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
HP-UX 11.31	✗	✗	✗	✗

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
AIX 6.1	✓	✗	✓	✗
AIX 7.1	✗	✗	✓	✓ 6.0.1.0+
AIX 7.2	✗	✗	✓	✓

Supported SQL Server Versions, Operating Systems, and Backup Software

 Delphix only supports 64-bit operating systems.

Windows Server

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
• Delphix only supports 64-bit operating systems.

Key:

Color	Supported?
	Yes
	Not Applicable

Windows Server

Supported OS Version	Supported DBMS Version		
	SQL Server 2012	SQL Server 2014	SQL Server 2016
Windows Server 2012			
Windows Server 2012 R2			
Windows Server 2016			
Windows Server 2019			
Windows Server 2022			

Tested Browser and Operating Systems

This topic describes the Web browsers and operating systems that have been tested for use with the Delphix Management application.

The following table lists the supported browsers and operating systems.

OS	Browsers
Windows 11	Microsoft Edge, Firefox, Chrome
Windows 10	Microsoft Edge, Firefox, Chrome
Windows 8.1	Microsoft Edge, Firefox, Chrome
Windows 7	Microsoft Edge, Firefox, Chrome
Mac OS X	Microsoft Edge, Firefox, Chrome

 Delphix supports the latest version of Microsoft Edge, Firefox, and Chrome as well as the immediate prior version.

Fixed Issues (6.0.x)

Release 6.0.13.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-80818	Fixed an issue where some tables in MDS are missing indexes when upgrading, causing the <code>libc</code> upgrade to necessitate a PostgreSQL re-index.

Release 6.0.13.0 Changes

Security Fixes

Bug Number	Bug Introduced in	Description	Security Bulletin
DLPX-79789	5.3.0.0	Arbitrary Code Execution May Be Performed by Engine System Administrators	TB096

Management Server Fixes

Bug Number	Description
DLPX-78689	The issue with the confusing vPDB error on a snapshot after resetlogs of a linked CDB is now fixed.
DLPX-41671	You can now update the Oracle cluster home through the user interface.
DLPX-43467	When dSource is an Oracle standby in RTA mode, LogSync was raising the <code>fault.oracle.linkedsources.log.conflict</code> error and getting disabled on its own. This issue is now fixed.
DLPX-50309	Users can now change <code>logSyncInterval</code> for Oracle dSources.

Bug Number	Description
DLPX-59757	The count for masking jobs fetched from the Masking Engine is now configurable. By default, it is set to 500.
DLPX-67604	The manual recovery of a database after V2P from a snapshot of dSource was failing with an error. This issue is now fixed.
DLPX-68684	The self-signed certificate is now compliant with the requirements for trusted certificates in MacOS 10.15.
DLPX-74613	Oracle VDB migration check needs to be done against the Oracle target host instead of the source. Furthermore, the Error and Action plan provided should include the target hostname.
DLPX-75467	The CLI now returns correct and descriptive error messages when executing unauthorized requests on uninitialized engines.
DLPX-75646	To diagnose BEQ connection failure, this release adds MD5 checksums for ojdbc*.jar for Oracle release versions up to currently supported release version.
DLPX-75878	The issue with the JDBC connection string for an Oracle vPDB not getting updated after an IP address change is now fixed.
DLPX-75989	The issue with the failure of environment discovery of an Oracle Cluster with a NullPointerException error is now fixed.
DLPX-76956	Previously, the Oracle JDBC test connection with the wrong password was increasing the <code>LCOUNT</code> value by more than 1. This issue is now fixed.
DLPX-77140	This release now speeds up metadata that is sent during replication when Extended retention is involved.
DLPX-77231	Previously, when source discontinuity on the dSource was followed by resync on the livesource, one or more livesource workers were failing to start. This prevented livesource status from getting updated and the first snapshot from being taken after resync. This issue is now fixed.
DLPX-77600	This release fixes NPE when Linking dSource with missing backup and unreachable nodes.
DLPX-77880	This release improves scalability for engines with an extremely large number of snapshots that were causing them to run out of memory.

Bug Number	Description
DLPX-78015	Previously, V2P export with absolute data files was failing with an internal error. This issue is now fixed.
DLPX-78174	Insecure DES is no longer supported for SNMPv3.
DLPX-78420	This release adds V2P support for Windows server 2022 host machines.
DLPX-78473	This release improves load times for Datasets and Dataset Performance pages for engines with a large number of datasets and containers.
DLPX-78488	Previously, when switching from the backup server to ASE, dump history was not working for dSources configured to use the remote backup server. This issue is now fixed.
DLPX-78594	This release fixes an issue with disabled VDBs not being able to undo a refresh.
DLPX-78688	TLS 1.0 and TLS 1.1 ciphers are no longer available. Any system that is only configured with TLS 1.0 or TLS 1.1 ciphers is switched to use the default cipher set.
DLPX-78693	Invalid sync parameters will not cause DE server unavailability.
DLPX-78696	Switching sync strategy from source sync strategy type to sourceless strategy type is not allowed.
DLPX-79126	VDBs can now be automatically started with the tunables if stopped intermittently because Windows fail to write on the mount (Msg 9001).
DLPX-79292	This release eases restrictions on taking a snapshot of PDBs with encrypted UNDO tablespaces.
DLPX-79344	Previously, Snapsync of a standby PDB in mount mode was failing with the <code>ORA-01109: database not open</code> error message. This issue is now fixed.
DLPX-79422	Previously, clicking on a Replication Profile was resulting in an unknown error message dialog. This issue is now fixed.

Bug Number	Description
DLPX-79789	Under certain conditions, arbitrary code execution may be performed by sysadmins.
DLPX-79808	This release fixes failures if the VDB name is more than 68 characters.

Release 6.0.12.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-79151	The issue with remote syslog configurations preventing engine upgrades or virtualization service restarts is now fixed.

Release 6.0.12.0 Changes

Log4j Updates

Based on detailed testing and analysis, all the currently supported products are not susceptible to known log4j vulnerabilities. Please refer to [TB095 Technical Bulletin](#) for more information. All instances of log4j in currently supported Delphix products are updated to **log4j 2.17.1** as of this release.

Delphix keeps you updated on the latest developments and keeps releasing hotfixes, procedures, and workarounds for such critical vulnerabilities. For more information on how Delphix supports our product and customers in such cases, see [Delphix Product Security](#).

For more information, refer to the following pages:

- [TB095 log4j Vulnerabilities](#)
- [Uninstalling the Delphix Connector Service from the Target Database Servers](#)
- [Delphix Product Lifecycle Policies](#)
- [Product Security](#)

Management Server Fixes

Bug Number	Description
DLPX-23068	Validation for target database parameter 'DB_FILES' for single-tenant databases for the following operations: provisioning, refresh, rewind, and converting a dSource to LiveSource is now added. Furthermore, specific error messages for handling ORA-00059 errors are added.

Bug Number	Description
DLPX-44 544	The issue with the SnapSync of an Oracle standby dSource in Real-Time Apply mode calculating the snapshot's timestamp incorrectly is now fixed. This issue was resulting in ORA-01194 or ORA-01152 errors when provisioning to a timestamp after the snapshot.
DLPX-56 691	The issue with the data files of a VDB getting unmounted when provisioning, refresh, or rollback job is canceled manually is now fixed.
DLPX-57 971	The issue with the latest snapshot of a LiveSource taking a long time to show the SCN/timestamp range on its card in the GUI is now fixed.
DLPX-60 320	UI now allows the selection of older dataset repositories (downgrade) in addition to selecting newer ones (upgrade).
DLPX-67 069	The issue with the stopped Oracle VDB monitoring by the environment monitor that resulted in connection errors flooding the debug log is now fixed.
DLPX-68 132	The issue with the “Copy query to clipboard” SQL copy functionality in the “Managed Source Data” is now fixed to have correct apostrophe characters.
DLPX-72 123	The issue with the failure of detaching or deleting an Oracle dSource operation on RAC environments (This issue was occurring due to failure of deletion of RMAN backups on RAC and the operation needs to be retried with a force option) is now fixed.
DLPX-72 779	The UI showing enabled or disabled for cluster nodes now uses a grid table. The Enabled column contains a checkmark to show whether the cluster is enabled or disabled. Users are able to select or unselect the checkmark to enable or disable a cluster.
DLPX-73 975	Critical storage faults should not be ignorable nor manually resolvable
DLPX-74 862	This release fixes an issue where the RESUME operation failed without any error thrown to the user on dSources where ENFORCE was still in progress. The fix will make sure that even if RESUME fails, it throws a DUE to the user with suggested actions to resolve the issue.
DLPX-75 763	The issue with the failure of refreshing a VDB provisioned as an empty vfiles since there is no parent container to refresh from is now fixed.
DLPX-76 266	The issue with the VDB Disable operation that resulted in an error message while connectivity with the host machine can't be established is now fixed.

Bug Number	Description
DLPX-77 123	You can now run Upgrade Verify when another upgrade is in progress.
DLPX-77 347	This release fixes the difference in time shown for MSSQL snapshot based on different database authentication methods.
DLPX-77 638	The issue with the failure of the End Entity Certificate expiration fault is now fixed.
DLPX-77 664	The issue with the failure of the Oracle SnapSync with an error message, "RMAN-06183: datafile or datafile copy (file number) larger than MAXSETSIZE" if a datafile resized in the middle of SnapSync is now fixed.
DLPX-77 913	The issue with the faults table missing data if there was more than one page of faults is now fixed.
DLPX-77 925	The issue with the unsupported Windows release error message is now fixed.
DLPX-78 113	MSSQL VDB database size will now be refreshed periodically based on environment_monitor.dynamic_attributes_check_period tunable.
DLPX-78 183	The issue with the MSSQL validated sync schedule not getting updated without successful backup restoration is now fixed.
DLPX-78 244	The issue with the failure of a few operations on self-service containers due to incorrect entries corresponding to the Oracle log metadata on the Delphix engine is now fixed.
DLPX-78 258	The issue with the input bug that retained cleared out DB credentials in the dSource linking wizard is now fixed.
DLPX-78 263	The issue with the failure of a SnapSync of an Oracle standby dSource in Real-Time Apply mode with an error message, "exception.oracle.snل.linkedsource.current_scn.invalid" if the rate of change in the database is low is now fixed.
DLPX-78 265	Offline Oracle bystander PDBs data files can now be optimized leading to improved provision performance.
DLPX-78 309	The issue with the CLI being unable to log in to system users when the main virtualization service is down is now fixed.

Bug Number	Description
DLPX-78334	The issue with a large number of missing Oracle archive logs causing an error while viewing dataset is now fixed.
DLPX-78392	Hosts running Windows Server 2022 can now be added as Source and Target environments to the Delphix Engine.
DLPX-78522	SSLv3, TLS 1.0, and TLS 1.1 are no longer configurations options for HTTPS. Any system configured only with these removed options will be automatically set to use TLS 1.2.
DLPX-78791	This release upgrades log4j from 1.2.17 to the latest 2.x in Windows Connector.
DLPX-78938	This release upgrades log4j in virtualization to 2.17.1.

Release 6.0.11.0 Changes

Security Fixes

Bug Number	Bug Introduced in	Description	Security Bulletin
DLPX-77921	6.0.8.0	Arbitrary Code Execution by Delphix System Administrators may be Performed on Virtualization and Masking Engines	TB094

Management Server Fixes

Bug Number	Description
DLPX-53019	The issue with the missing redo alert raised during the environment monitor check has now been resolved.
DLPX-59299	Discovery and monitoring rely on "Connected" in sqlplus output, which may not be the case if NLS_LANG is set to another language (e.g. Japanese). Downstream operations, like linking or provisioning, may then fail due to missing user privileges.
DLPX-59662	The issue with copy-Only Backups failure with Virtual Service Accounts has now been resolved.

Bug Number	Description
DLPX-62 706	The issue with the Hostchecker not properly checking /home/delphix permissions has now been resolved.
DLPX-64 082	The issue with Oracle Provisioning scripts having hard-coded timeout issues has now been resolved.
DLPX-65 729	Added retry functionality to the 'read backup files' operation during a validated sync to an account for an unstable environment.
DLPX-69 778	SAML response is not logged on successful SSO login.
DLPX-72 043	The issue where LiteSpeed <code>xp_restore_headeronly</code> stored procedure failure message are displayed when validated sync is active for dSources with LiteSpeed backup has now been resolved.
DLPX-72 220	A UI issue that occurred while updating the vault when only the private key is changed has now been resolved.
DLPX-72 225	Admin user created from management UI is no longer showing as 'non-admin' type.
DLPX-72 237	The 'Verify Credentials' button from the DSP Throughput test page is now removed.
DLPX-72 369	The dependency on a parent snapshot relying on the latest snapshot is now removed if a parent snapshot does not exist during the VDB enable operation.
DLPX-72 778	Oracle dSource attach operation with changed DB ID using 'Force' option is now allowed.
DLPX-74 555	Updated the "no Delphix connector" message while provisioning a Windows source environment.
DLPX-74 676	Oracle LiveSource LogSync should only catalog valid archive log files.
DLPX-74 851	In the Add Environment GUI, the mouseover information for "Set Delphix Session Protocol Options (DSP)" has been corrected.

Bug Number	Description
DLPX-74 896	The race condition issue when running Oracle VDB refresh and dSource snap sync resulting in incorrect internal references has now been resolved.
DLPX-75 335	Added a product name and product version for the Delphix Connector executable so this information can be available before installation.
DLPX-75 500	For ag cluster nodes, if the refresh fails due to timezone discovery failure, don't delete the nodes from MDS as it doesn't mean we had an issue with the nodes.
DLPX-75 952	Database configs will be replicated only if the associated VDB is replicated.
DLPX-75 995	The issue causing environment 'Add' or 'Refresh' to fail when PowerShell Transcription is enabled has now been resolved.
DLPX-76 244	The issue where TCP fallback connection to database stops responding if the Oracle database instance is down has now been resolved.
DLPX-76 290	Databases of UNKNOWN cdb type are now included in the attachment of a non-PDB container.
DLPX-76 731	Added Delphix support for <code>WALLET_ROOT</code> and <code>TDE_CONFIGURATION</code> parameters to manage wallets in 19c instead of sqlnet.ora.
DLPX-76 759	Added "Response" to faults along with other details when logged in the Admin App.
DLPX-76 777	Remove orphaned Oracle logs resulting from archive log fetch timeouts.
DLPX-76 793	Added execution timeout for execution of <code>UpdateFileACL.ps1</code> .
DLPX-76 974	The issue where a user was unable to change 'from address' of SMTP server to noreply@delphix.com in the GUI has now been resolved.
DLPX-77 112	The issue where an Oracle VDB cannot be provisioned between different minor versions if the Source is on higher RU has now been resolved.

Bug Number	Description
DLPX-77 284	The issue where after a hotfix was removed due to a successful upgrade, the system would still indicate the hotfix was installed post-upgrade has now been resolved.
DLPX-77 345	The issue where provisioning a vVDB fails with <code>java.lang.OutOfMemoryError</code> when sqlplus is used to rename the datafiles has now been resolved.
DLPX-77 405	Replicated password vaults will no longer be visible in the UI.
DLPX-77 676	The issue where provisioning a vPDB from a PDB dSource fails with "ORA-65114: space usage in container is too high" if PDB <code>max_size/max_pdb_storage</code> is configured has now been resolved.
DLPX-77 708	The issue where a refresh/disable/destroy of a VDB using NFSv3 could cause loss of access to other VDBs that were using NFSv3 has now been resolved.
DLPX-77 844	The issue where V2P operations from a VDB snapshot would result in the deletion of any production datafiles that exist on specified V2P target directory has now been resolved.
DLPX-77 904	Removed 'Factory Reset' for Delphix Engines that are Data Vaults, as the operation is disabled for those engines.
DLPX-77 912	The issue that can cause a VDB stop, refresh, or rollback to fail with an internal error has been resolved.
ORB-246 5	Removed the requirement that SAML SSO email addresses must match case-sensitively for SSO logins.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-77577	Increased nvme I/O timeout to prevent storage issues in EC2.

Release 6.0.10.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-77 467	Loading the setup app dashboard (as sysadmin) was rendering a server error popup with instruction to contact Delphix Support. This 6.0.10.0 error has been known to impair the ability to configure web proxy, PhoneHome, SMTP servers, and other connectivity settings via the GUI. It has now been resolved.

Release 6.0.10.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-18 438	The issue with provisioning to the latest available time that resulted in generating the <code>exception.oracle.target.point.not_provisionable</code> exception has now been resolved.
DLPX-35 480	Previously, static routes were being added using the add command. Now, the same can be added using the create command.
DLPX-57 516	The issue with the failure of management service to start after configuring some abbreviated timezones from a picklist in server setup or sysadmin CLI has now been resolved.
DLPX-58 133	Previously, the Oracle SnapSync operation was resulting in a warning for BCT usage on editions that do not provide it. This issue has now been resolved.
DLPX-58 675	The issue with the deletion of the last snapshot on timeflow by Retention during a failed Oracle DB_SYNC operation has now been resolved.
DLPX-63 003	The issue with memory being exhausted while reading too many snapshots from MDS has now been resolved.
DLPX-63 347	If the staging source has the "Use as Staging" flag set as off, the user was seeing a specific exception while trying to enable a linked dSource to point in the direction of what needs to be done. Any compatibility failure will now have a specific exception.
DLPX-63 601	Previously, querying the following operation "backupset table" and "whether a database is part of AG or not" was resulting in deadlocks and lock timeouts errors. We have now added retries to resolve the issue.

Bug Number	Description
DLPX-64 369	The issue with throwing <code>fault.oracle.linkedsource.incomplete.tempfile</code> for physical standby has now been resolved.
DLPX-65 949	The issue with misleading status in the progress bar while taking a copy-only backup has now been resolved.
DLPX-69 831	Previously, the Oracle dSource SnapSync operation was not displaying a clear failure message if a dSource <code>db_unique_name</code> is changed. This issue has now been resolved.
DLPX-70 317	The issue where the reaper thread cannot unmount a timeflow that was recently marked not in use, because the container associated with it (a VDB) got deleted, has now been resolved.
DLPX-71 018	The issue with UI displaying only the suffix of the device name used by Hyper-V has now been resolved. UI now displays a unique device name for storage in Hyper-V.
DLPX-71 292	The issue with the allowance of incremental SnapSync after LogSync throws <code>fault.oracle.linkedsource.log.conflict</code> has now been resolved.
DLPX-71 639	The NFSv4 is now set as the default option when mounting datasets from OS platforms that support it.
DLPX-71 769	The need to set permissions of <code>\$ORACLE_HOME/dbs</code> subdirectories using STARTUP SPFILE syntax is now removed.
DLPX-72 011	The issue with the CLI network setup not configuring the first network interface when multiple interfaces exist has now been resolved.
DLPX-72 186	The issue with CDB log file retention working incorrectly if a PDB has multiple time flows pointing to the same CDB timeflow has now been resolved.
DLPX-72 432	Support bundle enhancements have been made in this release. A timestamp is recorded for each sample in the log, making it easier to determine the time for each sample.
DLPX-72 956	The issue with disabling the Oracle LogSync after running the validated sync job has now been resolved.
DLPX-73 575	The timezone monitoring is now added for the Windows hosts.

Bug Number	Description
DLPX-73 590	You can now refresh a VDB whose parent dataset is in a different group without needing authorization on the parent or its group.
DLPX-73 800	The issue with the failure of olsnodes when run as a non-Oracle user has now been resolved.
DLPX-74 504	The issue with throwing a new DUE and NotFoundException when ojdbc libs cannot be read has now been resolved.
DLPX-74 945	UI now displays a detailed error message for transaction log-chain break fault.
DLPX-74 992	The issue with the failure of SnapSync operation when Database incarnation reset-logs end time is changed from "2021-03-13 22:03:07.0" to "2021-03-13 21:03:07.0" for virtual pluggable database "<database>" has now been resolved.
DLPX-75 389	The issue with the recording of the insufficient details by Logsync when dbid change was detected has now been resolved.
DLPX-75 517	The issue with the failure of the Oracle vPDB provisioning with the "ORA-00959: tablespace 'TEMP' does not exist" error has now been resolved.
DLPX-75 721	The issue with the failure of an Environment discovery with the "DelphixFatalException: Unknown Oracle Database status: REFRESHING" error has now been resolved.
DLPX-75 737	The issue with saving unnecessary logs by Retention if bookmark falls exactly on a snapshot end SCN or snapshot end timestamp has now been resolved.
DLPX-75 897	Previously, failure to start I/O services after the upgrade operation was resulting in a stack restart loop. This issue has now been resolved.
DLPX-75 951	An error signaled during VDB SnapSync when a data file is physically removed during processing by an external cause has now been resolved.
DLPX-76 288	The NFS latency for workloads involving a lot of parallel I/O (e.g. Oracle VDBs with concurrent accesses to many data files) is now improved.
DLPX-76 388	Previously, entering key pairs directly into hook environment variables, as opposed to via a vault or as passwords were resulting in an internal error. This issue has now been resolved.

Bug Number	Description
DLPX-76406	The issue with NFS-based VDBs becoming unresponsive has now been resolved.
DLPX-76447	The issue with the V2P Functionality to customize target directory structure for exporting database files to separate file systems not working as documented has now been resolved.
DLPX-76613	The issue with the unnecessary accumulation of heap when validated sync is active for dSources using an environment user that eventually can cause out of memory issues has now been resolved.
DLPX-76690	The issue with the removal of extraneous Oracle data files while creating snapshots has now been resolved. The extraneous Oracle data files are now removed during the provisioning operation.
DLPX-76692	For V2P operation, we now use the unbuffered copy method for better performance.
DLPX-76718	The issue with the creation of the extraneous Self-Service branch segments during the replication operation has now been resolved.
DLPX-76760	Environment clusters will now show faults from their child nodes.
DLPX-76802	The issue where the starting verification container enables IP forwarding on host has been resolved.
DLPX-76891	CRON expression labels now ask for Quartz format on the user interface.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-76619	The NFS reliance on DNS to prevent VDB unresponsiveness related to DNS unreliability is now reduced.
DLPX-76203	The NFS latency for engines with many Oracle dNFS clients is now improved.
DLPX-76119	The issue with the Delphix Engine crashing or becoming unresponsive when canceling a replication job has now been resolved.

Bug Number	Description
DLPX-76991	Optimized in-memory cache eviction by making minor improvements to I/O performance.

Release 6.0.9.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-28435	MS SQL instances PatchLevel will be displayed in preference to the version on the UI.
DLPX-40005	Improved the error message that is displayed when a JDBC connection cannot be established or the Oracle database Instance is unavailable.
DLPX-48080	Oracle Home Check may generate spurious faults if Oracle Home entry does not exist in /etc/oratab.
DLPX-55951	Attempting to provision a plugin-based VDB onto an incompatible OS (Windows to Linux or vice versa) is possible in the UI, and would fail with a crash requiring a restart. Now an informative error message is shown after the attempt is made.
DLPX-59613	Fixed creation of retention policies workers on replica objects after failover.
DLPX-64307	Environment refresh should ignore cluster discovery for Oracle VDBs.
DLPX-66191	Fixed the side-effects of the native Windows "Recent Files" behavior, when large numbers of PowerShell operations are being run concurrently.
DLPX-67537	Domain administrators can now create, view, and edit the alert profiles of other domain users.
DLPX-69605	Poor error message when selecting Timeflow range.
DLPX-70502	On detaching a dSource, delete backup server entry from MDS and related ones if the backup server is unused.

Bug Number	Description
DLPX-71002	In case we have null values coming for recovery_model from msdb.dbo.backupset table, the user will see a generic exception for manual sync and a fault for validated sync.
DLPX-71908	Users will now see a warning when they remove any object from the replication specification list.
DLPX-72012	Prevents the same IP address from being configured on more than one network interface.
DLPX-72411	Environment names for Windows and Oracle Clusters are once again editable by users.
DLPX-72609	When we do MSSQL standalone environment discovery, the user will see a warning for databases attached to AG that are present and will not be discovered unless cluster environment discovery is selected.
DLPX-72695	Improved Oracle SnapSync performance by eliminating unnecessary calls to getTotalHoleBlocks.
DLPX-73409	Duplicate listener entry gets generated in MDS if Oracle listener is manually started with a non-uppercase name.
DLPX-73586	Fixed a display error of some snapshot names in the command line interface which showed references instead of actual names.
DLPX-73720	Provisioning an Oracle vPDB fails with "ORA-65149: PDB name conflicts with existing service name in the CDB or the PDB" if the PDB and CDB names are the same.
DLPX-74050	Environment names for Windows and Oracle Clusters are once again editable by users.
DLPX-74078	The Target Directory path is combined with other directories such as Data Directory, Archive Directory, Temp Directory, etc to build the full path for data files, archive logs, temp files, etc. As long as the combined paths are valid the V2P job proceeds.
DLPX-74197	Fixes a misleading warning about insufficient space on a replication target.

Bug Number	Description
DLPX-74201	Snapshots created as a result of refresh or rewind operations will now be labeled as just "Snapshot" to avoid confusion. Users are advised to look at the Timeflow markers to know when the Timeflow operation was performed.
DLPX-74367	Delphix Engine repeatedly reports "Failed to parse logfile".
DLPX-74377	Improved diagnostics information for the case when Delphix Engine fails to connect to the Windows host.
DLPX-74387	fix an issue that causes the management service to crash under heavy CPU load.
DLPX-74398	Added handling of dangling nodes during Windows cluster environment add and refresh operations.
DLPX-74486	Delphix OS users cannot provision 12.2 TDE vPDB due to directory permissions in the default wallet location.
DLPX-74495	Enabled more logging in Delphix connector logs for timeouts.
DLPX-74681	During RAC vPDB provision, Oracle 19.9 target CDB crashes with ORA-00600 [krccf_chunk] when BCT is enabled.
DLPX-74806	Improved error message displayed when the storage device initialization fails unexpectedly.
DLPX-74860	Provisioning the 2nd generation VDB fails if the dSource has imported read-only transportable tablespaces fails.
DLPX-74975	Allow adding invalid or unreachable paths as a shared backup location for dSources.
DLPX-75026	Invalid JDBC connections are not purged from the connection pool when the home is changed.
DLPX-75134	Improved performance for the Environment Databases pages when there are a lot of databases.

Bug Number	Description
DLPX-75363	Update exception description is seen when the SnapSync fails for ASE encrypted database.
DLPX-75401	local listener set to null if oracle.lsnr.protocol_registration_order is quoted.
DLPX-75506	Fix a bug that can cause Oracle RAC VDBs to fail with stale NFSv3 mounts if NFSv4 is also enabled.
DLPX-75532	Insufficient heap memory settings on AIX cause connector and SnapSync to hang or crash.
DLPX-75663	FIPS compliant algorithms will be used while merging the old and new toolkit directories during environment refresh.
DLPX-75716	Delphix may remove Oracle VDB temp tablespaces during Snapsync.
DLPX-75735	Fixed creation of retention policies workers on replica objects after failover.
DLPX-75834	Rearranged Syslog configuration dialog inputs to avoid confusion and have a more consistent user experience.
DLPX-75858	BEQ processes can hold on to file descriptors leading to hook scripts hanging after upgrading to 6.0.7.0.
DLPX-76018	Remove hardcoded 5-minute timeout for doDropPDBKeepDatafiles.sh.
DLPX-76140	TDE SnapSync should ignore WARNING plugin violations.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-71980	Fix a bug that was causing the Delphix Engine storage pool to fail to import on boot under certain circumstances.

Release 6.0.8.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-75804	Some Delphix operations may fail if mount and umount commands, on staging or target hosts, are setup to run as sudo and if sudo rules prohibit these commands from running with unrecognized options. The issue is fixed now after removing "-v" added in 6.0.8.0.

Release 6.0.8.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-49694	Skip parsing of DBCC when code 0 is not present in the DBCC page output. In case the DBCC page has code 0 but not dbid, use bdbid (present in the buffer section).
DLPX-68764	VDB SnapShot does not progress if ASE database devices are not on Delphix storage, environment monitoring raises a fault. Subsequent VDB operations like enable, disable, start, stop, delete, snapshot, or refresh will fail.
DLPX-70793	Delphix Engine should not allow linking Oracle DB with null <code>db_unique_name</code> .
DLPX-71300	For newer ASE versions (\geq 15.7 SP138 and 16.0 SP02 PL05 and ASE 16.0 SP03), do not run DBCC PAGE anymore, as it was an identifier for DBCC CHECKALLOC that is already not run.
DLPX-71471	Error message asking user to manually perform disable/enable operation or correctly rename the target database back will be displayed during Start VDB, if VDB does not exist.
DLPX-71687	Provide a mechanism to enable VDBs up to filesystem mount point.
DLPX-71875	Fixed a bug that results in a memory reservation not being represented correctly in the Delphix API.
DLPX-72046	Deletion of vPDB in a vCDB shows this warning, "Encountered an error while shutting down and cleaning up Oracle files."

Bug Number	Description
DLPX-7 2209	Downloading a support bundle is not supported at the same time that an upload of an upgrade image has been initiated by the same Delphix user.
DLPX-7 2319	Fixed an issue where some error dialogs would freeze in Internet Explorer 11.
DLPX-7 2705	Connection timeout when deleting remote shipper script can cause a timeout in LogSync client.
DLPX-7 2757	ASE sync using Dump History fails for large dump history files.
DLPX-7 2780	Timezone is set incorrectly for snapshots of Solaris 10 dSources and VDBs.
DLPX-7 2904	Storage capacity now includes usage from all file system objects, not just snapshots.
DLPX-7 3143	Fixed an issue where the support bundle dialog showed a loading spinner intermittently while jobs were running.
DLPX-7 3354	Traverse all shared backup locations while syncing, even if some of the paths are invalid or not reachable.
DLPX-7 3489	Fixed bug where adding or editing a parameter using the UI VDB Config Template "Text" tab was truncating the parameter's value.
DLPX-7 3602	Incorrect mount options used when a single instance RAC is linked as a standalone single instance.
DLPX-7 3607	Added paging for days with large numbers of snapshots to prevent slowdown.
DLPX-7 3623	Fixed an out-of-memory condition that occurs in SSH tunneling for encrypted log-syncs when storage latencies are high.
DLPX-7 3627	The help text on upgrade replication warnings have been updated to avoid confusion between Ignored and Resolved.

Bug Number	Description
DLPX-7 3668	Fixed Missing security headers.
DLPX-7 3669	Cross-site request forgery (CSRF) issue in management UI.
DLPX-7 3727	Fixed an issue where the faults table was unable to navigate to other pages.
DLPX-7 3797	Fixed VDB refresh failures due to SQL Server Error 924 after setting VDB to single user mode.
DLPX-7 4025	Implemented logic to retry offline database along with a drop database to overcome deadlocks while off-lining or dropping the database.
DLPX-7 4029	VMware Hot-Add memory is not immediately reflected in the system API.
DLPX-7 4057	Fixed a typo in "Download Support Bundle" UI component where the word "suport" was missing a "p".
DLPX-7 4254	Ownership of files inside VDB now matches new owner when VDB owner is changed.
DLPX-7 4298	Fixed an issue where the user could not upload a keystore with a blank keystore passcode.
DLPX-7 4362	Fixes an issue with namespace deletion when the replication receive jobs have been cleaned up.
DLPX-7 4442	VDB Enable with attemptStart=false will now mount the datasets so that VDB can be started.
DLPX-7 4457	Cluster discovery for Oracle RAC clusters are partially failing on Solaris 10.
DLPX-7 4529	Fixed a bug so that an upgrade completes even when jobs fail.

Bug Number	Description
DLPX-74542	Fixed a bug so that upgrade completion is properly handled after kernel upgrades.
DLPX-74645	Delphix Engine uses the uptime command to keep track of a target host reboot and auto start VDBs on the host. In some cases, the output of this command is not what is expected and causes unintended restart of a stopped VDB. This issue is now fixed.
DLPX-74656	Oracle errors during doCreateSPFile.sh are not captured.
DLPX-74704	Fixed a bug where the Dataset scroll does not extend to the bottom of a dataset list, thus truncating the status of the last dataset in the expanded group.
DLPX-74883	Prevent support bundle collection from cancelling replication.
DLPX-74911	Talaria TCP fallback fault may be misconstrued if an Oracle RAC node is down.
DLPX-74997	Prevent granting replicated roles to users.
DLPX-75083	Post upgrade cleanup task may become unresponsive while attempting a migration from 5.3.x to 6.0.x due to several threads stuck in WAITING state.
DLPX-75095	Provisioning an Oracle VDB fails if change-archivelog-mode.sh takes longer than 5 minutes.
DLPX-75134	Improved performance of the Environment Databases page under certain conditions.
DLPX-75188	Fixed "out of memory" issue when processing a large number of objects on the Target engine.
DLPX-75204	Addressed a performance issue on the Target engine when receiving large number of replicated objects.
DLPX-75208	Issue with redacted snapshot names in support bundles have been resolved.

Bug Number	Description
DLPX-75416	Fixed a replication issue when there are sources with TLS enabled.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-72065	Fixed a bug that can cause a Windows iSCSI initiator to fail connect to the Delphix Engine.
DLPX-72681	Console Delphix status screen shows a Python stack trace if the system is configured with a static IP address.
DLPX-73423	Console Delphix status screen shows a Python stack trace if the system has no default route.
DLPX-74216	Fixed an issue that causes management service failures in low memory situations.
DLPX-74622	Fixed a bug that can cause a replication job to fail with an internal error.
DLPX-75089	Fixed a bug that can cause NFSv3 clients to lose locks during upgrade verification.
DLPX-75524	Fixed a bug that can lead to Oracle data corruption when running VDBs on Oracle 19c with dNFS.

Release 6.0.7.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-39006	LogSync failed with "Cannot read archived log due to failure of log shipping script".
DLPX-39245	Fixed a bug that caused the management service to become inaccessible if the storage pool ran out of space.
DLPX-59155	Provisioning a VDB or vPDB failed with unclear error message 'A database with the name "xxx" already exists'.

Bug Number	Description
DLPX-60 317	Fixed Out of Memory issue when replicating a large number of objects.
DLPX-60 947	Replica VDBs will be updated when performing a point-in-time restore.
DLPX-62 805	vPDB provision did not raise an error when a non-provisionable target point-in-time was provided.
DLPX-62 969	Fixed Out of Memory issue when receiving large number of replicated objects.
DLPX-64 600	Skipped connecting to ASE dSources during SnapSync policy runs as it is not applicable for them, hence prevent recurrent faults that the policy throws for connectivity issues.
DLPX-67 363	Maximum identify provider authentication time age can be customized for single sign-on.
DLPX-67 607	Fix to make Snapsync throw exception if manifest file is missing or of 0 bytes instead of internal error with null pointer exception.
DLPX-67 767	Fixed a bug that caused the upgrade to hang, while waiting for running jobs to finish.
DLPX-70 821	Allow the entity id for SAML single sign-on to be a URL for compatibility with Azure AD.
DLPX-71 783	doRenameDatafiles cleanup of extra files fails due to file permissions mismatch.
DLPX-72 010	Fixed an issue that prevents changing the default gateway using the network setup CLI.
DLPX-72 075	Maximum SAML response time skew can be customized for single sign-on.
DLPX-72 191	Oracle privilege discovery not performed for all homes if an invalid home exists.
DLPX-72 351	When a user tries to change credentials for a dSource, validating the credentials before updating them. In case of invalid credentials, showing user an error message about it.

Bug Number	Description
DLPX-72 545	Initial ORA-65294 error not reported to user when vPDB provision fails due to compatible parameter mismatch.
DLPX-72 652	Fix and issue that prevents use of the NFSv4 on some versions of SUSE Linux targets.
DLPX-72 698	Patching Oracle 19C vCDB leads to ORA-25153 as described in 2285159.1.
DLPX-72 807	Fixed issue with SQL Server 2014 dSources with filestreams where sync failed in merging filestream directories due to long path names.
DLPX-72 882	Datasets hooks script editor properly displays multiline scripts instead of as one long line on non-Chrome browsers.
DLPX-72 916	Empty string in SNMPv3 USM username creation no longer throws fatal error.
DLPX-73 048	Non-sys user credentials for Oracle sources cannot use password vault.
DLPX-73 108	Fix a bug that prevents the API from displaying the correct number of CPUs or amount of memory assigned to a Delphix Engine after a hot-add operation.
DLPX-73 201	Fix an issue that prevents the configuration of additional NICs on Azure Delphix Engines.
DLPX-73 202	Fix a bug that can cause a VDB to fail to mount while other VDBs are being stopped.
DLPX-73 424	Fix a bug that prevents the sysadmin from deleting a default route.
DLPX-73 527	SnapSync job fails with 'internal error during execution' due to ONS/FanManager errors.
DLPX-73 528	Fixed a bug that prevented accessing SDD specs from CLI.
DLPX-73 611	Kerberos ticket expiration date parsing is incorrect after migration from Illumos to Linux.

Bug Number	Description
DLPX-73742	Provisioning an Oracle TDE-enabled vPDB fails with the error "ORA-28367: wallet does not exist" if the TDE wallet for the target linked CDB is stored on ASM storage.
DLPX-73765	Fix a file descriptor leak that causes the management service to crash over time.
DLPX-73789	Auxiliary CDB instance uses dSource keystore location if WALLET_ROOT is configured in dSource.
DLPX-74030	CDB database password may be leaked as part of environment monitor checks that launch sqlplus command on the source or target host.
DLPX-74043	Delphix OS user cannot provision TDE-enabled vPDB due to directory permissions in the default wallet location.
DLPX-74044	Delphix OS user cannot provision TDE-enabled vPDB in Delphix-writable keystore location due to directory permissions.
DLPX-74119	Drop database fails if default database is set to any other than master.
DLPX-74164	Sync fails with db.aselddb.source.dump_history.incomplete_stripes after dump history file is purged and Use dump history is enabled for the dSource.
DLPX-74233	During failover of a namespace, if there is a collision between an environment in the namespace with one on the target engine, the namespace environment will get renamed if its host does not match that of the environment on the target.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-73390	Improve replication receive throughput.
DLPX-73393	Improve write performance under extreme disk fragmentation.
DLPX-73280	Improve write performance under extreme disk fragmentation.

Release 6.0.6.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-73848	Fixed an issue that can cause the management service to fail to start after upgrade on systems that have had SNMP enabled.
DLPX-73859	Fixed a file descriptor leak triggered by faults and alerts that can cause the management service to fail.

Release 6.0.6.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-47065	VDB recovery failed when files other than archive logs were detected by Oracle.
DLPX-47493	Fixed the bug where VDB directory under the DelphixConnector directory was not being removed from the target host on MSSQL VDB deletion.
DLPX-48046	Added sorting parameter to network test APIs.
DLPX-61405	Replication may send more data than expected if masking involves dropping large DBF files.
DLPX-61525	The height of the storage configuration list was limited to show 3 disks at a time. It will now dynamically grow with the number of disks.
DLPX-63603	Increased connector timeout from 10 minutes to 30 minutes to avoid unnecessary faults due to timeout during Validated Sync operation.
DLPX-67368	Delphix Engine hostname change is now immediately reflected in Splunk events.
DLPX-67593	Fixed an issue that caused the management service to remain offline following an out-of-space condition.

DLPX-68 531	Introduced better handling of UniversalConnectionPoolException errors during SnapSync.
DLPX-69 759	Oracle environment discovery failed due to an unhandled exception occurring at insert into dlp_x_faults.
DLPX-69 852	Fixed a bug that caused network configuration problems when removing and adding additional NICs.
DLPX-70 426	Redaction of usernames took forever on tables with millions of entries.
DLPX-70 583	move-to-asm.sh fails if timing is set in glogin.sql.
DLPX-70 638	Removed Failed Actions section of Actions sidebar, in favor of manually dismissing from Running Actions and falling to Finished Actions section.
DLPX-70 653	Removal of all instances in a RAC VDB should not be allowed.
DLPX-70 808	Fixed issue related to the creation of empty DisableBroker.sql on the Windows machine in case DisableBroker.sql execution fails in the first attempt.
DLPX-70 896	Added more detailed error message for when the Delphix Engine fails to push a script to Windows host.
DLPX-70 919	Fixed an issue that causes job progress to not update in Self-Service.
DLPX-70 928	Fixed a bug that results in a Delphix Engine remaining powered on following a shutdown from the user interface.
DLPX-71 093	For AG databases, a full backup is not required even recovery fork guid changed but the LSN chain didn't break because of transactional log backups.
DLPX-71 097	Unable to ignore snl.bct.needed warnings if Block Change Tracking is legitimately disabled on an Oracle dSource.
DLPX-71 153	Recovery of PDB should fail if the database is down after offlining datafiles.
DLPX-71 370	While deleting initiator in Windows environment deletion operation, delete all the views as well for that initiator.

DLPX-71 685	VDB is auto disabled if the hook fails.
DLPX-71 865	Reduced the size of support bundles.
DLPX-71 961	When a PDB is selected for replication, its CDB and all other PDBs in the parent CDB get automatically selected for replication. Going forward, in the above scenario, while the CDB will get selected, its other PDBs will no longer get selected.
DLPX-72 031	Fixed VDB refresh operations failures due to 'DB STARTUP' background process spid greater than 50.
DLPX-72 066	Migrate VDB verifies against the old configuration, rather than new.
DLPX-72 083	Fix an issue that causes a fully-qualified hostname to be changed on upgrade from 5.3 to 6.0.
DLPX-72 131	Added namespace support for HashiCorp password vaults.
DLPX-72 265	doCreateTempfiles.sh.template exits with code 0 on failure.
DLPX-72 340	Incomplete recovery not detected during provisioning.
DLPX-72 386	Unlock Solaris x86 Solaris -> Linux x86 provisioning.
DLPX-72 452	For clusters with long hostnames, vPDB sync fails with exception.oracle.accessor.instances.missing.
DLPX-72 495	Fixed a bug that prevents the application from coming up after an upgrade
DLPX-72 686	Delphix no longer logs environment variables in logs on connected hosts since this could leak sensitive information such as passwords that are sometimes stored as environment variables on database hosts such as for the ASE database.
DLPX-72 730	Fixed a Snapsync performance issue.

DLPX-72790	SnapSync job fails with 'internal error during execution' due to ORA-01652.
DLPX-72862	The scenario which was causing the null pointer has been fixed now.
DLPX-73300	Validation of connection to a container for PDBs should allow connecting to CDB\$ROOT.
DLPX-73311	Added platform detection for ESX 7.0u1.
DLPX-73449	Replication of policies between two engines, in a loop, could lead to OOM exceptions.

Delphix OS Fixes

Bug Number	Description
DLPX-72990	Addressed a minor CVSS 5.9 security issue with no known attack vectors.
DLPX-73067	Fix for CVE-2020-10753.
DLPX-73069	Fix for CVE-2020-12059.
DLPX-73070	Fix for CVE-2020-1760.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-71924	Fixed a bug that causes support bundle collection to fail.
DLPX-72918	Fixed a system crash that can happen when replicating a masked VDB using SDD.
DLPX-73147	Fixed a bug that can cause a replication source to crash if it had run replication while running on 5.0.

Release 6.0.5.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-23 360	The unistallation of the Delphix Connector installer should succeed even if one of the component connector services doesn't exist.
DLPX-69 155	Reduced the time taken to generate support bundles in some cases.
DLPX-67 316	Recreating a controlfile against an Oracle source may yield misleading error during snapshot.
DLPX-70 766	JDBC driver updated to resolve intermittent JDBC connection failures due to JDBC SSL bug.
DLPX-70 785	Options passed to VDB mounts on target AIX hosts did not include read and write size values. This fix adds the rsize and wsize parameters to mount command depending on the maximum values host is configured to support.
DLPX-70 741	Enabling Validated Sync while SAP ASE SnapSync job is running leaves staging database unrecoverable.
DLPX-69 865	Fixed a bug that causes a network interface to become unconfigured if its MAC address changes.
DLPX-71 233	If LogSync is suspended when performing SnapSync of a standby database in real-time apply, SnapSync attempts to backup the archived logs which can cause SnapSync to become unresponsive.
DLPX-69 800	UEM/Hostchecker directory ownership checks fail on HPUX environment with long usernames.
DLPX-69 807	Provided mechanism for the user to bypass corrupted/incomplete jdbc libraries.
DLPX-66 585	Bundle ID "fault.environment.configuration.file.owner" reports insufficient host address.
DLPX-65 739	createDelphixDBUser.sh fails when "@" used in the password.

DLPX-70 973	SAP ASE database provisioning fails if the source database has holes in log fragments.
DLPX-71 532	Improved error handling for Oracle memory configuration errors.
DLPX-62 987	Allowed assigning privileges over replicated objects through the UI.
DLPX-71 593	TIMEFLOW_REPAIR incorrectly skips a log because of "wrong database".
DLPX-71 751	Added NFSv4 support on AIX for Oracle and SAP ASE.
DLPX-71 736	Dynamically disable RPC services if NFSv3 is no longer in use.
DLPX-71 772	Network DSP Test between versions 5.3 and (6.0.3, 6.0.4) is fixed.
DLPX-71 513	Replicate non-data objects like delphix engine users, authorizations, roles, permissions, policies and DB config templates.
DLPX-71 305	Unable to load dummy recovery database dump due to SAP ASE error 15728.
DLPX-71 172	Enabling SAP ASE dump history causes IllegalStateException in getDumpListFromLastRestoreDateAndFiles due to timestamp mismatch because of TZ.
DLPX-71 178	SAP ASE internal error raised when dump history file is purged using sp_dump_history.
DLPX-65 101	Fixed a race condition between a DB_DELETE job and the Oracle retention policy worker for the same container that can lead to a deadlock between the job and the worker.
DLPX-71 918	Fixed an issue that causes the Delphix Engine UUID to change upon rebooting in IBM Cloud.
DLPX-71 141	Fixed an issue where upgrading an Oracle dSource or changing the environment user in a linked Oracle dSource fails with the error "SOURCE_UPGRADE job for xxx failed due to an internal error during execution."

DLPX-71 611	Updated UI time zone library to IANA 2020a.
DLPX-70 349	Fixed a memory leak during incremental replication.
DLPX-72 038	Fixed an issue that prevents 5.3.x - 6.0.x upgrade if a static route exists that goes over a DHCP interface.
DLPX-72 148	Fixed a bug of always order hooks alphabetically rather than the running order set by users.
DLPX-71 971	Allowed Enable/Disable of VDB if its current Timeflow has at least one snapshot.
DLPX-72 115	Changed the Time Point field on a VDB back to reflecting the point on the parent the VDB was created from, but displayed in the timezone of the parent.
DLPX-71 995	6.0.4.0 can no longer interact with 5.3.x remote Masking Engines.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-70675	Fixed a bug that causes the system to become unresponsive after expanding multiple storage devices.

Release 6.0.4.2 Changes

Management Server Fixes

Bug Number	Description
DLPX-72 155	Fixed an issue that can render a Delphix Engine unbootable if a reboot occurs after upgrade verification but before the upgrade is applied.
DLPX-71 141	Fixed an issue where upgrading an Oracle dSource or changing the environment user in a linked Oracle dSource fails with the error "SOURCE_UPGRADE job for xxx failed due to an internal error during execution."

Release 6.0.4.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-71930	Fix a bug that causes feature flags to be disabled when upgrading to 6.0.4.0.

Release 6.0.4.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-68173	Resolved an issue where temporary database backup/device files created for cleaning up the target database were not being deleted.
DLPX-68773	The management stack runs out of memory as the environment monitor does not purge stale objects.
DLPX-69573	Allow linking an Oracle PDB with a lowercase name.
DLPX-69634	Allow provisioning an Oracle PDB with a lowercase name.
DLPX-69962	After detaching a PDB, perform unplug/plug, and attach again, if disabled is performed before SnapSync, the PDB can no longer be enabled.
DLPX-66045	Prevent Self-Service Container branches getting into an unusable state by blocking deleting the last segment of branches.
DLPX-7037	Snapsync performs an unnecessary checkpoint.
DLPX-68277	Users will see the detailed error message upon connection failure to Delphix connector during OS user validation and there will also be a "More" button with an error message which will open an error popup with all error details.
DLPX-70288	On the "Add Environment" screen when OS user validation will get fail, they will see the "More" button along with the error message. When the user clicks the button, an error popup opens with all details of the error and suggested action.

Bug Number	Description
DLPX-70 832	NFSv4 support for appdata sources running on AIX.
DLPX-68 495	Fixed GUI reporting conflict information when creating a Retention Policy.
DLPX-70 788	Added Environment User field for MSSQL sources in Datasets -> Configuration -> Source tab -> Staging Environment section.
DLPX-58 047	Fixed bug where the sort sequence was incorrect. Fixed in Hook Operation Templates.
DLPX-67 931	Provision against VPDB after create/drop a new tablespace failed with exception.oracle.targetscripts.rename.datafiles.
DLPX-59 910	Comps.xml associated with Oracle Homes are marked as unparseable if they are longer than 65535 characters.
DLPX-55 476	CLI provisioning fails when the mount point provided includes quotes around the path.
DLPX-71 168	Changed type to text and spaced "Secret Key" and "Username Key".
ORB-328 5	Support using api.delphix.com as a proxy for verifying the Cloud Agent binary's code signature certificate.
DLPX-71 006	Allow provisioning across patch versions for Oracle versions on or after 18.X.
DLPX-71 334	Migrate NTP configuration when upgrading between 5.3 and 6.0.
ORB-328 6	Communication with Central management servers is now routed through the web proxy when one is configured for the Engine.
ORB-311 7	Summary: Increase an action's failure message size to 256 characters so users can view large failure messages.

Release 6.0.3.1 Changes

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-71339	Fixed an issue that can cause the Virtualization Management service to become inaccessible when the system memory became highly fragmented.

Release 6.0.3.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-63192	More details will be displayed in the error message in case DB_SYNC fail due to missing SELECT permission on database 'msdb'.
DLPX-67708	Removed unnecessary Source Continuity source-archive file system.
DLPX-66878	A meaningful error message will be displayed in case the user is missing 'VIEW ANY DEFINITION' and 'VIEW SERVER STATE' permissions on AG Instance and linking is performed.
DLPX-68668	Fixed the issue when Environment discovery fails with Internal Error when Oracle DB instance name is > 15 characters.
DLPX-68539	Added support for Read-Only Oracle homes.
DLPX-62027	Fixed a bug that causes factory reset to fail when there are provisioned VDBs.
DLPX-67830	Eliminated a virtualization management service crash caused by egregious use of memory by environment monitoring.
DLPX-69067	Enabled NFSv4 support for older RedHat NFS clients.
DLPX-68491	Fixed an issue when SCAN Listener is not discovered for Oracle 19c Cluster Environment.

DLPX-68 931	Improved replication throughput by parallelizing data streams.
DLPX-69 350	Fixed an issue where in some cases a VDB's Time Point would not appear.
DLPX-62 602	Prevented a full snapsync after detecting an incarnation change of and reverting to a previous incarnation of an Oracle database.
DLPX-69 579	Resolved the issue of intermittent failure of DB_SYNC for source database full backups containing in-memory tables which were caused due to improper merging of filestream folders.
DLPX-69 104	Fixed an issue when Environment Monitor task monitors replicated entities could lead to Out of Memory.
DLPX-58 561	Increased online Redo Log size when using VDB Provisioning defaults from 50mb to 1024mb.
DLPX-68 831	Storage is removed even when the drop database fails, causing ASE error 823.
DLPX-68 323	Linking will not fail in case a slash is used as the path delimiter on the source database.
DLPX-69 561	Allowed NoLogging Diagnosis to be shown and edited for Oracle CDBs.
DLPX-69 625	Fixed an issue that causes the CLI to hang when deleting an object.
DLPX-65 357	Source Environment selection in Attach dSource dialog is now alphabetically sorted.
DLPX-65 215	Fixed an issue where Hotfixes aren't listed until after management service restarts.
DLPX-57 988	The increased timeout of doShutdownOracleInstance.sh script from 20 seconds to 10 minutes.
DLPX-70 018	Resolved the issue where during validated sync, fault "fault.mssql.source.next.backup.missing" was caused due to backupsets with similar first and last lsn.
DLPX-66 671	dSource selection in dSource Linking Wizard is now alphabetically sorted.

DLPX-65 723	MSSQL server cluster address is now editable through the Environments GUI.
DLPX-69 514	Gracefully handle accelerated networking on Azure.
DLPX-68 942	Implemented retries with some time delay in case of a failure while switching database user mode.
DLPX-56 626	Some orcl_log_info entries have a very large and incorrect end_scn (281474976710655).
DLPX-67 579	Deleted users' actions should be included in the action/audit log API results.
DLPX-69 863	Enhanced instruction text relating to SSH when editing environment users.
DLPX-70 081	Removed excessive debug logging for DSP connections which results in fast rollover of debug logs.
DLPX-66 203	CLI / API calls to refresh/rewind vCDB directly should be disallowed.
DLPX-67 194	RHEL 7.6 connector log shows Unidentifiable version string: RedHatEnterpriseServer 7.6.
DLPX-66 754	When VDB is disabled, environment configuration can now be edited in the UI.
DLPX-62 095	A wrong certificate is identified as an issuer of a self-signed certificate in rare cases.
DLPX-69 243	Do not require an issuer to be present or keep the full chain intact on Truststore operations.
DLPX-59 331	Permit non-CA certs in user Truststore.
DLPX-60 779	Changed error message when there are no compatible installations on provisioning.
DLPX-69 518	Provisioning failures due to BitLocker encryption will be identified and a proper error message will be displayed.

DLPX-64 797	Fixing memory leak in hk2 library.
DLPX-70 039	Password vault migration nullifies ASE linked source dump credentials.
DLPX-70 089	Protection against a variant of billion laughs attack (XML entity expansion).
DLPX-64 207	Added API support to revert from static to DHCP DNS settings.
DLPX-68 857	Faults reported for Oracle Home missing where the Central Inventory does not show this Oracle Home present.
DLPX-64 435	Exclusively specifying 'required' parameters to discover Oracle cluster via CLI results in an exception.
DLPX-69 604	Alerts & Faults are reported for hosts in a namespace that can cause Out of Memory issue.
DLPX-39 882	Prevented cloning of Tiimeflow storage for Oracle source continuity.
DLPX-67 425	Resolved an issue when validated sync (with full/diff) restored multiple backupsets and a restore failed with a SQL server transient issue after a source continuity reset event resulting in a state where no operations could be performed on the dSource.
DLPX-70 433	"DLPX_EXECUTE_SQL_CLEANUP_RETRY" will also print nested SQL error messages in case of command failure.
DLPX-68 582	Customers now have access to an API to display the Engine License information.
DLPX-61 335	Displayed in confirmation dialog the name of the user being deleted.
DLPX-70 639	Resolved output buffer issue while identifying BitLocker encryption during provisioning on Win19.
DLPX-66 259	Updated messages on the upgrade page when the operation fails.

DLPX-70782	Bumped up connector version for NET 4x installer as shipped OpenJDK version had been upgraded.
------------	--

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-68995	Improved performance of dataset deletion.
DLPX-68997 DLPX-68999	Improved single connection replication throughput.
DLPX-70697 DLPX-70703	Addressed an issue that causes long periods of I/O unresponsiveness.
DLPX-69953	Fixed a bug that can cause a Windows iSCSI initiator to fail to connect to the Delphix Engine.
DLPX-70512	Fix a hang in the I/O subsystem that can cause the Delphix Engine to become unresponsive.

Release 6.0.2.1 Changes

Management Server Fixes

Bug Number	Description
DLPX-70065	Provisioning a VDB from a dSource or another VDB will fail if the following conditions are met: <ul style="list-style-type: none"> Delphix Engine has at least one dSource and a VDB created using a Python plugin prior to the upgrade Delphix Engine was upgraded to 6.0.2 Provisioning was attempted from the UI after the completion of the upgrade
DLPX-69350	Fixed an issue that the time point attribute of a VDB is not shown.

Delphix OS Fixes (Activated after optional Reboot)

Bug Number	Description
DLPX-69864	Fixed an issue that causes MSSQL operations to hang after the reception of an iSCSI LUN reset.

Release 6.0.2.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-62806	Fixed an issue where provision against PDB after unplug/replug against the same linked PDB fails with exception.oracle.targetscripts.controlfile.create.
DLPX-67567	Oracle Source Continuity creates an unnecessary source-archive file system on zfs.
DLPX-27807	LogSync may fall behind when connected to an Oracle physical standby database in Real-Time Apply mode.
DLPX-68385	Customer provided Oracle Java missing in the search path for Java on hosts.
DLPX-62782	Reducing the number of nodes for RAC VDB and VDB in NOARCHIVELOG mode may result in ORA-00258 errors during VDB enable operation.
DLPX-62738	Better error message when plugins are uploaded out of sequence.
DLPX-68722	The product now recognizes VMware with BIOS date of 12/12/2018 as VMware 6.7.0u2.
DLPX-68579	SnapSync of Oracle 19c DB with encrypted tablespace fails with fatal exception "Block header 91 is not empty".
DLPX-68689	Fixed the issue where a huge number of error messages from ASE caused OutOfMemory Error.
DLPX-68957	Always On AG discovery will not fail in a multi-subnet environment.
DLPX-63088	Can now recover multiple Self-Service containers at the same time.
DLPX-47977	Improved handling of snapshot standby.
DLPX-64125	SnapSync failed with exception.oracle.dsource.sync.no_hosts.rac on RAC clusters with very long hostnames.

Bug Number	Description
DLPX-62584	PDB enable failed after migration if mountBase has a trailing slash.
DLPX-68657	Virtualization can now fetch jobs from Masking engines configured with HTTP redirection.
DLPX-69121	It is no longer mandatory to have at least one enabled system administrator with local credentials.
DLPX-68167	Fixed an issue where too many requests were being sent for Faults from the Datasets pages.
DLPX-69082	Large stderr produced by failed rsync jobs are truncated to prevent Java OutOfMemory errors.
DLPX-58600	Datasets filter updated so that all items within a group that matches the filter string are displayed, even if the items contained in the group do not match the filter string.
DLPX-65896	VDB deletion failed due to the inability to delete LogSync worker.
DLPX-57903	Improved diagnosability of PDB discovery issues.
DLPX-68878	Fixed issue where start/stop buttons were not being displayed in the RAC instances configuration table.
DLPX-69271	Enabled replication smart failover by default.
DLPX-66715	The user-visible name for Oracle cluster objects is being replaced with the Oracle cluster name. For Windows clusters, the user-visible name is being replaced with the cluster address.
DLPX-68929	Changed default replication settings for better out of the box performance.
DLPX-68930	Improved replication throughput when sending multiple timeflows.
DLPX-69245	Fixed a memory leak that occurs when experiencing connectivity errors.

Bug Number	Description
DLPX-69377	At least one non LDAP system user should be enabled when the LDAP server is being disabled.
DLPX-68575	LDAP principal fields were not being redacted in phone-home bundles.
DLPX-68528	Self Service Recover operation failed due to missing Timeflow.

Delphix OS Fixes

Bug Number	Description
DLPX-66808	Re-introduced console splash screen with IP address and service states.

Release 6.0.1.1 Changes

Delphix OS Fixes (Optional on Reboot)

Bug Number	Description
DLPX-69203	Improved synchronous write performance over iSCSI.
DLPX-69167	Improved SQL Server data ingestion performance by leveraging asynchronous writes on underlying storage.
DLPX-69298	Eliminated possible data corruption on SQL server and vFiles over iSCSI that can occur when a Delphix Engine reboots.

Release 6.0.1.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-60689	For SAP ASE, instead of using the DBCC CHECKALLOC command to fix DBID mismatch issue, the MOUNT command with FIXDBID and ALLOW_DBID_MISMATCH clauses will be used, to improve performance.

Bug Number	Description
DLPX-658 31	VDB snapshots need to clean unneeded ZFS datafiles.
DLPX-639 49	Improved boot time after 5.3 to 6.0 migration by optimizing metadata indexing.
DLPX-662 61	Upgrades to 6.0.0.0 will only be supported from a release greater than or equal to 5.3.6.0.
DLPX-664 86	Snapshot of a linked database can end up with extra datafiles that do not belong to the database which might cause VDB on VDB provision to fail during rename of datafiles.
DLPX-665 58	Cluster environment discovery was incomplete if the host locale was not English.
DLPX-668 04	DB_LINK using incorrect user when RAC node also configured as a standalone environment.
DLPX-667 68	vPDB save state lead to rollback or child provisioning failures.
DLPX-668 23	Unable to link database with CL8MSWIN1251 charset.
DLPX-645 38	Fixed a bug causing the timezone selector to only be visible when manually setting the time.
DLPX-668 09	Removed the Windows Diagnostics Files and Directories on successful Diagnostics upload.
DLPX-672 79	Provision failed when the source was from a RAC Oracle Standard Edition database and the target was Oracle Standard Edition.
DLPX-674 51	Fixed an issue that sporadically caused replication to fail with an internal error.
DLPX-674 54	Delphix Engine should select the highest version ojdbc driver available at ORACLE_HOME/ojdbc/lib.
DLPX-660 77	Ensures child worker threads are gracefully exited when parent linked source sync job has completed/terminated.

Bug Number	Description
DLPX-459 83	MSSQL Validated sync will resume when storage usage falls below the threshold if storage threshold enforcement failed in the past.
DLPX-675 60	Fixed an issue where MT provision may result in ORA-02058 due to un-purged or inflight 2PC transactions on dSource.
DLPX-675 94	Old timeflows and snapshots are not getting removed by snapshot retention.
LX-2020	Report the correct amount of memory allocated to EC2 Nitro instances.
DLPX-674 13	Fixed an issue where VDB point in time provisioning might fail if Oracle database environment is configured in a non-English locale.
DLPX-676 84	PDB provisioning failed if the source had shutdown triggers.
DLPX-675 75	Fixed failure during point in time 'Virtual to Physical' provisioning.
DLPX-676 68	After setting the database online give it some extra time to startup completely, before doing any further operation on it.
DLPX-677 59	Redact sensitive information from phone-home data.
DLPX-646 38	Validated sync stops working if Delphix cannot connect to the backup server.
DLPX-655 59	Even when the staging instance is down, attempt counter to detect backup files keeps on increasing and eventually, it stops detecting backups.
DLPX-565 37	When a target host is used by a large number of dSources for staging or has a large number of objects, the performance of Delphix operations like validated sync, refresh, rewind, etc can be slow due to Powershell processes being serialized.
DLPX-678 94	Removing cluster resource without removing its dependency can result in cluster failure. So, added retryer logic while fetching the resource dependencies (Get-ClusterResourceDependency) and ultimately fail the operation after all the retries.

Bug Number	Description
DLPX-678 13	Unsupported SQL server backup type gets picked while validated sync and the operation fails while looking for the backup. So, introduced a tunable filter to automatically skip SQL backups taken by backup software not supported by yet Delphix.
DLPX-679 25	Added env host connectivity toolkit support for SLES on Power9.
DLPX-679 34	Retries to fetch image identifiers during Netbackup restore if there is a mismatch between MSDB and Netbackup Master.
DLPX-676 55	Fixed an issue where retention enforcement can generate user-visible errors while attempting to delete snapshots with dependencies after PDB migration to new CDB.
LX-1944	EBS NVMe devices can now be used in Delphix Engines.
DLPX-680 22	Fixed an issue where hostchecker 'Check Oracle DB Instance' fails on HPUX and AIX.
DLPX-681 24	PDBs with lower/mixed case names will not enable after an upgrade.
DLPX-681 26	Fixed a bug that limits the number of disks that can be added in GCP.
DLPX-674 21	Update the primary db file names in a transaction with the Timeflow creation to make sure whenever a Timeflow is created successfully we have its primary file information.
DLPX-674 40	Skip VDBs having its current Timeflow as null from 'PrimaryDbFileAvailabilityCheck' as these VDBs doesn't undergo queisceing and are recoverable by refreshing them.
DLPX-618 18	Linking wizard - Target Environment step - Privileged Credentials authenticates on the selected target now.
DLPX-681 17	Some non-Admin users, lack all permissions, are unable to login to upgraded engine.
DLPX-672 90	A wrong version input by user while manually adding a SQL Server instance, created issues in provisioning VDBs. SQL Server version will now be auto-discovered for manually added instances on adding or refreshing the environment.

Bug Number	Description
DLPX-66238	Updated error message to let know user that non discovered CDBs are filtered out from the list when linking a detached source.
DLPX-68457	When a target host is used by a large number of dSources for staging or has a large number of objects, the performance of Delphix operations like validated sync, refresh, rewind, etc can be slow due to Powershell processes being serialized.
DLPX-68484	Fixed the issue where 'lstart' column value of sysusages table was beyond the range of Integer data type by taking the Long data type to store the lstart value.
DLPX-68500	Fixed an issue where the NTP service is not started following a reboot.
DLPX-68290	Support bundle generation can be time-consuming if the engine has a large number of snapshots to process.
DLPX-67792	Fixed issue in grids in which the selection checkbox was unclickable.
DLPX-67555	Provision vPDB/vCDB fails with ORA-45900 if the parameter enable_pluggable_database is omitted when specifying database parameters for new vCDB.

Delphix OS Fixes

Bug Number	Description
DLPX-67782	Engines running 5.3 on EC2 i3 can now be migrated to 6.0.
DLPX-67961	Fixed an issue that prevents ssh access after switching to a static IP address.
DLPX-65948	Fixed a bug that could cause replication jobs to fail with internal errors
DLPX-68025	Improved boot time after 5.3 to 6.0 upgrade by reducing the overhead of setting ZFS properties.
DLPX-67868	Fixed a bug that can cause the management service to run out of memory when disabling the Splunk integration.

Release 6.0.0.0 Changes

Management Server Fixes

Bug Number	Description
DLPX-27 433	The analytics GUI network graph shows newly added NIC information without requiring a management service restart.
DLPX-33 998	If you add a hook script via the CLI, newlines are removed erroneously.
DLPX-40 094	Correctly set the default type for the parameters to all operations in the CLI according to the container type.
DLPX-43 215	Exclude sybsecurity from the list of auto-discovered databases.
DLPX-48 712	Java 6 packages are no longer included in the product image.
DLPX-48 280	When a user is set with the Provisioner role the 'provision' button does not appear, meaning anyone set with this role only is unable to provision VDBs.
DLPX-53 996	The Delphix Engine does not provide instructions to browsers to avoid caching HTTP responses (pages).
DLPX-54 740	Ensure Windows mount points are always unmounted as part VDB refreshes to prevent future VDB refreshes from failing due to "ERROR_ASSIGN_MOUNTPATH: failed to assign mount path for disk <disk_number> at <mount_path>, error -2147024751".
DLPX-55 282	In environments where the vPDB has been provisioned using a Delphix provisioned virtual CDB, shutting down the virtual PDB causes it to get into an incorrect "Cannot monitor" state, this has now been fixed to show the correct "Stopped" state.
DLPX-55 598	Fixed an issue where vPDB refresh/rollback triggers spurious vCDB restart jobs, after vPDB+vCDB auto-restart.
DLPX-55 829	Validated Sync can fail when monitoring ASE backup servers started by using the \$DSLISEN environment variable instead of the "-S" argument. This can be worked around by accessing \$DSLISEN in the RUN_XXXX script and pass it down as -S.
DLPX-55 958	VDBs with no snapshots failed to re-enable after a Delphix Engine upgrade, this has now been fixed.

Bug Numer	Description
DLPX-57 454	Display underlying ssh error when environment host connections fail.
DLPX-58 519	Enable Oracle LiveSource when LiveSource is in RESYNC_NEEDED state currently re-start Oracle Redo Apply. Oracle Redo Apply should not be restarted in this state.
DLPX-58 760	Fixed a TCP port leak in the network throughput test feature.
DLPX-58 845	Provisioning vFiles to the same host using different OS Environment Users no longer fails.
DLPX-59 772	The API to list all snapshots consumes a significant amount of memory when there are more than 100,000 snapshots on the engine.
DLPX-60 356	Fixed an issue where Oracle remote listener registration fails if set to empty string.
DLPX-60 603	Network settings dialog now displays actual MTU value rather than a checkbox.
DLPX-60 907	Fixed an issue where the Environment Monitor on Redhat 6.9 and 6.10 might throw unidentified version errors.
DLPX-60 979	When user configures connection strings manually, these connect strings can end up connecting to incorrect PDBs/CDBs causing invalid snapshots. Verify that each connection to a PDB/CDB connects to the expected PDB/CDB.
DLPX-60 993	Delphix backups create controlfile records; in rare circumstances, these records can cause invalid snapshots. To avoid this problem, remove Delphix backups control file records when using SCN-based SnapSyncs once a SnapSync completes successfully.
DLPX-62 094	Allow certificates to expire after issuer certificate expiration.
DLPX-62 241	Reduce SSH connections by temporarily preserving and reusing existing Delphix<->host connections where possible.
DLPX-62 781	Spurious job event "DISCOVERED_TO_MANUAL_ORACLE_CLUSTER_NODES" no longer shows up for non-Oracle RAC environment refreshes.

Bug Number	Description
DLPX-62 892	In Oracle versions 18c and 19c, an Oracle bug can prevent the datafile headers from being updated for a standby database when managed recovery is running, resulting in failed SnapSync operations. Alert the user that an Oracle patch might be needed.
DLPX-62 962	Removed unneeded EMPTY_RENEGOTIATION cipher
DLPX-62 998	Fixed an issue where stale file mounts may be leftover when vPDB provision fails.
DLPX-63 469	Initial setup now fails if the system was not provisioned with enough storage.
DLPX-63 600	Network settings dialog now displays actual MTU value rather than a checkbox.
DLPX-64 641	Fixed an issue where the last snapshot of a vPDB Timeflow can be deleted after the vPDB has been disabled, thus leaving the vPDB in a state with no provisionable snapshots.
DLPX-64 711	Allow provisioning to complete when source CDB includes PDBs in a broken state.
DLPX-66 020	Provision should remove files present in datafile filesystem that are not part of the database when provisioning a VDB from a VDB.
DLPX-67 299	ASE environment discovery will not fail if there is a mismatch of "dataserver name argument" and value of "@@servername".

Delphix OS Fixes

Bug Number	Description
DLPX-57384	Fixed a system hang caused by a deadlock in ZFS.

Known Issues

Release 6.0.13.0

Key	Summary	Workaround
DLPX-38908	LogSync should automatically resolve faults for transient issues.	Manually resolve the faults generated when Oracle LogSync encounters transient issues.
DLPX-56320	While adding an Exadata source environment (RAC on Linux) to the Continuous Data Engine, the following error occurs: "The string must not be longer than 15 characters."	This workaround involves setting the Environment Name to a value longer than 15 characters after the environment discovery has completed. More information can be found in this Knowledge Base article.
DLPX-57078	Job cancel requests during provisioning are not processed until the end of the step.	Find and kill the Oracle <code>pmon</code> process for the instance that is being provisioned.
DLPX-64386	LogSync thread hangs when trying to remove temporary RMAN command file from source host toolkit directory.	Remove the <code>rm</code> alias for the Delphix OS user.
DLPX-76382	Force disable should succeed despite environmental problems.	No workaround.
DLPX-78412	SCM/Talaria failure reason should be communicated in fault.	No workaround.
DLPX-78986	Prevent DSP connections for disabled cluster nodes.	No workaround.
DLPX-79212	While restoring full backups of file stream type database on open staging DB and taking snapshots, snapshots are consuming full space instead of referring to each snapshot.	No workaround.
DLPX-79355	V2P export throws an error when suspended and resumed at the Opening Database step.	Cancel the V2P export job and start the V2P export operation all over again.

Key	Summary	Workaround
DLPX-79502	SnapSync fails if more than 1000 tempfiles exist in the whole CDB.	No workaround.
DLPX-79596	DB files are getting deleted for staging DB if staging push dSource is forced disabled after performing a restore outside Delphix and then enabled again.	No workaround.
DLPX-79833	You can no longer select performance metric resolution for fluentd configurations.	No workaround.

Release 6.0.12.0

Key	Summary	Workaround
DLPX-67604	Manually recovering a database after V2P from a snapshot of dSource fails with an error.	No workaround.
DLPX-69775	Updating Oracle credentials with an empty string throws an error when Simplified Connection Management is enabled.	Use the Delphix CLI to unset the user.
DLPX-75209	Network configuration is lost when changing EC2 instance type.	On Nitro-based instance types, use the virtual machine's virtual serial console and login to the sysadmin CLI, then add a DHCP address to the network interface.
DLPX-75878	The JDBC connection string for an Oracle vPDB does not get updated after a listener port change.	A second environment refresh will update the connection string for the vPDB.
DLPX-77231	When source discontinuity on the dSource is followed by resync on the livesource, one or more livesource workers may fail to start. This prevents livesource status from getting updated and the first snapshot from being taken after a resync.	No workaround.

Key	Summary	Workaround
DLPX-78689	Oracle vPDB snapshot job fails after doing reset logs on a linked CDB.	Disable the vPDB and re-enable it to clear out the snapshot job errors.
DLPX-78700	Oracle vPDB source enable jobs for Oracle 21c are taking more than 15 minutes to complete. This issue is only seen during the first Oracle vPDB enable operation on Oracle 21c.	No workaround.

Release 6.0.11.0

Key	Summary	Workaround
DLPX-44544	A SnapSync of an Oracle standby dSource in Real-Time Apply mode sometimes calculates the snapshot's timestamp incorrectly. This can cause ORA-01194 or ORA-01152 errors when provisioning to a timestamp after the snapshot.	To provision from a snapshot of an Oracle standby dSource in Real-Time Apply mode, provision by SCN instead of timestamp.
DLPX-57971	The latest snapshot of a LiveSource may take a long time to show the SCN/timestamp range on its card in the GUI.	No workaround.
DLPX-72123	Detaching or deleting an Oracle dSource may fail on RAC environments due to the failure of deletion of RMAN backups on RAC and the operation needs to be retried with the force option.	If the detach or delete operation on an Oracle dSource fails, retry the command using the force option.
DLPX-75209	Network configuration is lost when changing EC2 instance type.	On Nitro-based instance types, use the virtual machine's virtual serial console and login to the sysadmin CLI, then add a DHCP address to the network interface.

Key	Summary	Workaround
DLPX-7766 4	Oracle SnapSync fails with "RMAN-06183: datafile or datafile copy <datafile> (file number <fileno>) larger than MAXSETSIZE" if a datafile resized in the middle of SnapSync.	One of the following three workarounds can be applied: <ol style="list-style-type: none"> 1. Re-issue SnapSync. 2. Reduce the frequency of datafile resizes, or 3. Ensure datafile resize operations are not being performed while the SnapSync operation is in progress.
DLPX-7786 9	Unable to drop and recreate descending order (or any other functional) index as part of Oracle Connector masking/reidentification/tokenization jobs.	Oracle interprets descending order as functional indexes. There is no workaround. This is a known limitation of Drop Indexes for Oracle connectors that will be resolved in a future release.
DLPX-7801 5	V2P export with absolute data files is failing with an internal error.	No workaround.
DLPX-7826 3	A SnapSync of an Oracle standby dSource in Real-Time Apply mode sometimes fails with "exception.oracle.snل.linkedsourc.e.current_scn.invalid" if the rate of change in the database is low.	Force a log switch on all primary instances/nodes and then try another SnapSync.

Release 6.0.10.0

Key	Summary	Workaround
DLPX-7746 7	Loading the setup app dashboard (as sysadmin) renders a server error popup with instruction to contact Delphix Support. This message can be ignored. However, it is known that this error impairs the ability to configure web proxy, PhoneHome, and SMTP servers via the GUI.	These settings can still be configured via the CLI. This issue will be fixed in the next version release.
DLPX-6408 2	Oracle provisioning scripts have hard-coded timeouts.	Retry the provisioning operation to see if it succeeds. Otherwise, contact Delphix Support.

Key	Summary	Workaround
DLPX-72369	RAC migration for VDB with a deleted parent may fail with error "Cannot update RAC instances if virtual source 'xxx' has a deleted parent."	No workaround.
DLPX-74896	Race condition during refresh may result in incorrect MDS entry for parent snapshot.	No workaround.
DLPX-75148	DSP throughput tests do not work when from-version is < 6.0.6.0 and target-version is >= 6.0.6.0.	No workaround, however, the user will need to upgrade the source version.
DLPX-75215	Switching AWS instance types can leave the Delphix engine with no network configuration.	On Nitro-based instance types, use the virtual machine's virtual serial console and login to the sysadmin CLI, then add a DHCP address to the network interface.
DLPX-75995	On Windows 2019 Server, with KB4598230 cumulative update, adding or refreshing an environment fails when PowerShell transcription is enabled.	The user has to turn off the transcription if the staging is on Windows 2019 Server with recent updates.
DLPX-77231	When source discontinuity on the dSource is followed by resync on the Oracle LiveSource, one or more LiveSources workers may fail to start. This prevents the LiveSource status from getting updated and the first snapshot from being taken after a resync.	No workaround.

Release 6.0.9.0

Key	Summary	Workaround
DLPX-72186	CDB logfile retention works incorrectly if a PDB has multiple timeflows pointing to the same CDB timeflow.	No workaround.

Key	Summary	Workaround
DLPX-747 49	Oracle ENVIRONMENT_REFRESH_AND_DISCOVER job may fail with "The object OraclePDBConfig does not exist on the system".	Manually run an environment refresh after the ENVIRONMENT_REFRESH_AND_DISCOVER job fails.
DLPX-755 17	Provisioning an Oracle vPDB to a vCDB fails with "ORA-00959: tablespace 'TEMP' does not exist" if there is no temporary tablespace named TEMP within the production PDB\$SEED database.	One of the following two workarounds can be applied: <ol style="list-style-type: none"> 1. Re-attempt the vPDB provision to a linked CDB. The provision should succeed. 2. Manually create (or rename) the TEMP tablespace in the production PDB\$SEED database, take a new snapshot of the vPDB, and provision from that snapshot to the vCDB.
DLPX-757 37	Retention saves unnecessary logs if the bookmark falls exactly on a snapshot end SCN or snapshot end timestamp.	No workaround.
DLPX-759 95	On Windows 2019 Server, with KB4598230 cumulative update, adding or refreshing an environment fails when PowerShell transcription is enabled.	The user has to turn off the transcription if the staging is on Windows 2019 Server with recent updates.
DLPX-763 88	Entering key pairs directly into hook environment variables (a new feature in 6.0.9.0) results in an internal error.	Use password variables instead, or use a password vault. See Passing Credentials Securely to Hook Operations .

Release 6.0.8.0

Key	Summary	Workaround
DLPX-64307	Environment refresh should ignore cluster discovery for VDBs	<p>The following workaround should resolve the issue without Delphix support intervention.</p> <p>For RAC databases</p> <ol style="list-style-type: none"> 1. Add the instances to the clusterware configuration, ensuring you add all instances configured via Delphix. For example: <pre>srvctl add instance -db <Database Name> -instance <Instance Name> -node <Node Name></pre> 2. Refresh the environment. At this point, Delphix will add the instances back to the Delphix configuration, and you should be able to perform actions such as stop, start, disable, enable, etc. <p>From here, we recommend removing the clusterware configuration.</p> <ol style="list-style-type: none"> 3. Stop the VDB via the Delphix GUI or CLI. 4. Remove the instances from clusterware <pre>srvctl remove instance -db <Database Name> -instance <Instance Name></pre> 5. Remove the database from clusterware <pre>srvctl remove database -db <Database Name></pre> 6. Refresh the environment. 7. Start the VDB via the Delphix GUI or CLI. <p>For RACOne databases</p> <ol style="list-style-type: none"> 1. Stop the instance. <pre>sqlplus "/as sysdba" shutdown immediate</pre> 2. Ensure the RACOne configuration has the "instance prefix" configuration set. <pre>[oracle@mwtestx1 ~]\$ srvctl config database -d <vdb name> grep "Instance name prefix" Instance name prefix: VDB</pre> 3. If this is not set, then set it using <pre>srvctl modify database -db <vdb name> -instance <instance prefix></pre> 4. Set the pfile location to allow clusterware to start the instance <pre>srvctl modify database -db <vdb name> -spfile <path to pfile></pre>

Key	Summary	Workaround
		<p>The pfile is located in <Delphix Mount base>/<VDB Name>/datafile/ spfile.ora</p> <p>5. Start the instance using clusterware <code>srvctl start database -d vdb</code></p> <p>When querying the status of the database, it should now show a running instance. As an example:</p> <pre>[oracle@mwtestx1 ~]\$ srvctl status database -d VDB Instance VDB_1 is running on node mwtestx1 Online relocation: INACTIVE</pre> <p>6. Refresh the environment. At this point, Delphix will add the instances back to the Delphix configuration, and you should be able to perform actions such as stop, start, disable, enable, etc.</p> <p>From here, we recommend removing the clusterware configuration.</p> <p>7. Stop the VDB via the Delphix GUI or CLI. 8. Remove the database from clusterware <code>srvctl remove database -db <Database Name></code></p> <p>9. Remove the database from clusterware <code>srvctl remove database -db <Database Name></code></p> <p>10. Refresh the environment. 11. Start the VDB via the Delphix GUI or CLI.</p>
DLPX-721 86	CDB logfile retention works incorrectly if a PDB has multiple timeflows pointing to the same CDB timeflow	No workaround.
DLPX-734 09	Duplicate listener entry gets generated in MDS if Oracle listener is manually started with non-uppercase name	<p>Follow these steps:</p> <ol style="list-style-type: none"> 1. Restart the listener without a name parameter so it displays in uppercase. 2. Update VDBs to use uppercase LISTENER. 3. Refresh the environment. lowercase listener will be removed and LISTENER updated appropriately.

Key	Summary	Workaround
DLPX-747 49	Oracle ENVIRONMENT_REFRESH_AND_DISCOVER job may fail with "The object OraclePDBConfig does not exist on the system"	Manually run an environment refresh after the ENVIRONMENT_REFRESH_AND_DISCOVER job fails.
DLPX-748 60	Provisioning a 2nd generation VDB from a dSource with imported RO transportable tablespaces fails with ORA-19654	No workaround.
DLPX-748 82	Masking's SFTP client no longer compatible with SolarWinds and Goanyware SFTP servers	No workaround.
DLPX-755 17	Provisioning an Oracle vPDB to a vCDB fails with "ORA-00959: tablespace 'TEMP' does not exist" if there is no temporary tablespace named TEMP within the production PDB\$SEED database	One of the following two workarounds can be applied: <ol style="list-style-type: none"> 1. Re-attempt the vPDB provision to a linked CDB. The provision should succeed. 2. Manually create (or rename) the TEMP tablespace in the production PDB\$SEED database, take a new snapshot of the vPDB and provision from that snapshot to the vCDB.

Release 6.0.7.0

Key	Summary	Workaround
DLPX-7 4457	Cluster discovery for Oracle RAC partially fails on Solaris 10	Switch the default login shell for the Delphix OS user from /bin/sh to /bin/bash.
DLPX-7 4749	Oracle ENVIRONMENT_REFRESH_AND_DISCOVER job may fail with "The object OraclePDBConfig does not exist on the system"	Manually run an environment refresh from the Delphix UI if the ENVIRONMENT_REFRESH_AND_DISCOVER job fails with the mentioned error.

Key	Summary	Workaround
DLPX-76718	Time required to display the point-in-time pop-over on timelines for Self-Service Templates with replicated objects can degrade linearly with each replication.	No workaround but reducing replication frequency can reduce the impact.

Release 6.0.6.0

Key	Summary	Workaround
DLPX-73224	When provisioning from a non-multitenant source to a virtual pluggable database (vPDB), the post-plug hook script (dx-post-plug-hook.sh) should exit with the vPDB either closed or open unrestricted. If the vPDB is left open with restricted access subsequent Snapshots of the vPDB will fail with "oracle.ucp.UniversalConnectionPoolException" and "java.sql.SQLException: ORA-01035 ORACLE only available to users with RESTRICTED SESSION privilege" .	To prevent the issue: Ensure that the vPDB is either closed/ mounted or open unrestricted when exiting from dx-post-plug-hook.sh. After it has happened: close and reopen the vPDB.
DLPX-72749	Provisioning a TDE-enabled vPDB with system encrypted tablespaces fails with the error "ORA-28374: typed master key not found".	No workaround, contact Delphix Support (TBD)
DLPX-72655	Provisioning an Oracle TDE-enabled vPDB fails intermittently if the dSource is encrypted after linking	Enabling TDE on an existing non-encrypted dSource is not supported. Detach, rename the Delphix dSource name and re-attach it as a new TDE-enabled dSource before re-attempting the provisioning operation.

Key	Summary	Workaround
DLPX-73224	<p>When provisioning from a non-multitenant source to a virtual pluggable database (vPDB), the post-plug hook script (dx-post-plug-hook.sh) should exit with the vPDB either closed or open unrestricted. If the vPDB is left open with restricted access subsequent Snapshots of the vPDB will fail with "oracle.ucp.UniversalConnectionPoolException" and "java.sql.SQLException: ORA-01035 ORACLE only available to users with RESTRICTED SESSION privilege"</p>	<p>To prevent the issue: Ensure that the vPDB is either closed/ mounted or open unrestricted when exiting from dx-post-plug-hook.sh.</p> <p>After it has happened: close and reopen the vPDB.</p>
DLPX-73357	<p>Rewinding an Oracle TDE-enabled vPDB to the snapshot before vPDB migration may fail with <code>exception.oracle.tde.export.keys.failed</code> if all of the steps of the keystore merge procedure are not followed correctly</p>	<p>Verify that the CDB is restarted after merging the keys of the old target CDB into the new target CDB and re-attempt the rewind operation.</p> <p>If the rewind still fails, provision a new vPDB from the required snapshot instead of the rewind operation.</p>
DLPX-72181	<p>The "Restore Self-Service data container from the bookmark of sibling data container" operation may fail in an Oracle TDE environment with <code>exception.oracle.tde.export.keys.failed</code> if the keystore merge procedure is not followed correctly. In this case, the warning "<i>Refresh operation on the TDE-configured container <target container> from the Timeflow of another container <sibling container> requires merging of the TDE keystores.</i>" will be displayed in the Delphix UI events log</p>	<p>Verify that the procedure for merging the sibling keystores is followed correctly and re-attempt the self-service operation.</p>

Key	Summary	Workaround
DLPX-73742	Provisioning an Oracle TDE-enabled vPDB may fail with the error "ORA-28367: wallet does not exist" if the TDE wallet for the target linked CDB is stored on ASM storage	Storing the TDE wallet on ASM storage is currently unsupported. Modify <code>sqlnet.ora</code> to point to the keystore location outside of the ASM diskgroup and re-attempt the provisioning operation.
DLPX-73788	Provisioning or enabling an Oracle TDE-enabled vPDB fails when <code>\$ORACLE_BASE</code> is used in <code>sqlnet.ora</code>	Any environment variable referenced in <code>sqlnet.ora</code> must always be set in the environment for the Delphix OS user. Ensure that the environment variable <code>\$ORACLE_BASE</code> is set in the shell initialization file for the Delphix OS user and re-attempt the operation.
DLPX-73789	If <code>WALLET_ROOT</code> initialization parameter is configured on a TDE-enabled dSource PDB, provisioning may fail since the auxiliary CDB instance uses dSource keystore location.	If the provision fails, there are 2 workarounds: <ol style="list-style-type: none"> 1. Provide permissions to the Delphix OS user to create files in the location identified by <code>WALLET_ROOT</code> directory of the source keystore on the target host. 2. Modify the dSource database to not use <code>WALLET_ROOT</code> to identify the TDE keystore.

Release 6.0.2.0

Key	Summary	Workaround
DLPX-69638	Masking job created on engine 6.0.1.1 or prior is failing after the upgrade to version 6.0.2.0 or later	Masking jobs created in 6.0.1.x using a Hana JDBC driver will need to be updated to grant the following permission <code>{"java.io.FilePermission" "/", "read"}</code> in 6.0.2.0. All drivers created in and after 6.0.2.0 will be granted this permission by default.

Release 6.0.0.0

Key	Summary	Workaround
DLPX-60397	If a mapping algorithm is included in multiple jobs, only one job should be run at a time. If multiple jobs are run at the same time, then the mapping algorithm might contain multiple mappings to the same value or the jobs might deadlock.	Only run one job at a time.

Key	Summary	Workaround
DLPX-6 0947	Self-Service template with replica VDB is not updated with new Timeflow on incremental replication update	The latest replica VDB data can still be accessed by doing a Self-Service container Refresh, rather than a point-in-time restore from the template.
DLPX-6 1079	Certificate import validation may incorrectly reject a root CA certificate	Support must manually import the certificate into the truststore.
DLPX-6 1405	Masking operation should wait for zfs delete queue to drain	Replication may send more data than expected if masking involves dropping large DBF files.
DLPX-6 4493	V5 API /roles endpoint missing certain items	View and set these privileges through the GUI
DLPX-6 6155	Failed DSP engine test leads to multiple blocked client.jar processes on target hosts.	Restarting a Delphix Engine during a network throughput test is not recommended as it may lead to a system hang.
DLPX-6 6860	ADFS does not like NameIDPolicy sent by SSO app	Create an explicit rule in ADFS that transforms the {{emailAddress}} attribute into a {{nameid}}. The rule type must be "Transform an incoming claim". The incoming claim type must be "Email address" and the outgoing claim type "Name ID". The nameid format must be "Email address".
DLPX-6 6973	Date format is changed after importing the environment	Either (a) use the GUI import feature and then review the imported date formats for correctness or (b) use EngineSync to export/import jobs, which will not alter the date format.

Deprecated and End-of-Life Features

Release 6.0.12.0

End-of-Life Features

- **Internet Explorer 11 Support**
Internet Explorer 11 is no longer supported by Delphix. The users are requested to refer to the list of [supported browsers](#).
- **EBS 12.1 Support**
Starting 6.0.12.0 and beyond, EBS 12.1 is no longer supported by Delphix. The users are requested to [upgrade](#) to EBS 12.2 and start using this plugin.
- **SAP HANA 1.0 Support**
Starting 6.0.12.0 and beyond, SAP HANA 1.0 is no longer supported by Delphix. The users are requested to [upgrade](#) to HANA 2.0 and start using this plugin.

Release 6.0.11.0

Deprecated Features

- **Oracle 11.1**
Details of the Oracle database EOL can be found in the [Oracle Lifetime Support Policy](#) PDF.
- **TLS 1.0 and 1.1**
These versions of TLS are known to be vulnerable, enterprise use is heavily discouraged.

End-of-Life Features

- **Oracle 10g Support**
- **TLS 1.0 and 1.1**
These versions of TLS are known to be vulnerable, enterprise use is heavily discouraged.

Release 6.0.10.0

End-of-Life Features

- **EBS12.1**
EBS 12.1 remains deprecated with a planned EOL (and removal from the product) in December 2021. This is in line with Oracle's planned EOL for EBS 12.1 later this year. The users are encouraged to [upgrade](#) to Delphix's EBS 12.2.
- **Masking Ruleset Edit Options**
The Table Suffix, Add Column, Join Table, and List options were deprecated in the 6.0.3.0 release. These options have reached the EOL in the 6.0.10.0 release and have been completely removed from the product.
- **Internet Explorer 11 Support**
Internet Explorer 11 is planned for EOL in December 2021. This is in line with Microsoft's planned EOL for Internet Explorer 11 later this year.

Release 6.0.9.0

End-of-Life Features

- Delphix Reporting (Formerly, *Mission Control*) is EOL starting 6.0.9.0 (July 1st, 2021). The users interested in continuing their reporting workflows are recommended to discuss Data Control Tower with their account teams.

- Also, EBS 12.1 remains deprecated with a planned EOL (and removal from the product) in December 2021. This is in line with Oracle’s planned EOL for EBS 12.1 later this year. The users are encouraged to [upgrade](#) to Delphix’s EBS 12.2.

Release 6.0.8.0

End-of-Life Features

- SAP ASE 15.0.3
- Windows Server 2008 R2
- SQL Server 2008
- SQL Server 2008 R2

Release 6.0.7.0

Masking End-of-Life Features

- ESX 5.5 support will be end-of-life in 6.0.7
- Masking Connectors: Db2 LUW and zOS v9, Db2 LUW and zOS v10, SQL Server 2005, 2008, 2008 R2

Release 6.0.5.0

The following list of OS and DB versions are now past their respective support windows. Delphix will deprecate these versions to stay in line with Microsoft, Oracle, and SAP lifecycle policies:

- Windows 2008 R2: Microsoft End of Extended Support (EoES) January 4th, 2020
- SQL Server 2008/2008 R2: Microsoft End of Extended Support (EoES) July 9th, 2019
- Oracle 10: Oracle End of Extended Support (EoES) July 2013
- SAP ASE 15.0.3: SAP End-of-Life (EoL) March 31st, 2015

Release 6.0.4.0

End-of-Life Features

- **Masking:**
 - Job Scheduler - As of this release, Delphix has removed the Job Scheduler feature. The introduction of Masking’s REST API several releases ago allowed customers to schedule job executions using their preferred job scheduler. As a result, the integrated scheduler is seldom used.

Deprecated Features

- **Masking:**
 - FTP, SFTP, and mount upload for XML and Cobol formats - FTP/SFTP/Mount-based format import were the original modes for XML and Cobol files, since then, Delphix has added the ability to upload a format file, which is far simpler to set up. After the introduction of “upload”, we’ve seen a dramatic shift away from the legacy import modes in favor of the simplicity of “upload”.
 - Row Type Feature - Originally geared for limiting masking to subsets of rows within a column, this feature was seldomly used. Its functionality, if desired, can still be replicated via the Custom SQL feature.
 - Redundant Settings for ‘Edit Table’ under Rule Sets - Table Suffix, Add Column, Join Table, and List - These settings are redundant and can be replicated with the Custom SQL setting.
 - ‘HAVING’ clause from Masking API - Deprecating due to low use. This feature, if desired, can be replicated with Custom SQL.

Release 6.0.2.0

End-of-Life Features

- AWS i.3 (EOL) January 22, 2020
- Microsoft Azure GSx (EOL) May 7, 2020

Release 6.0.0.0

End-of-Life Features

- **Removed Virtualization Features:** A few features will be removed with the 6.0 release. If these features are in use, the upgrade will not proceed.
 - Cross Platform Provisioning (XPP)
 - AIX 6.1 Technology Levels 0-6 and AIX 7.1 Technology Levels 0-2 for Environments
 - Windows Server 2008 for Environments

 Windows Server 2008 will not have an upgrade blocker.

- **Removed Masking Features:** Please note that Excel files can still be masked by first converting them to one of Delphix's supported file types (CSV, etc). Also, XML CLOBs can be masked by extracting their values into a table (example - using extractValue in Oracle).
 - Native XML CLOB masking: After upgrade, columns masked as XML CLOBs will have the NULL SL algorithm assigned.
 - DB2 9.1, 9.5, and other 9.x versions of LUW & Z/OS
 - "Create target" job option: After upgrade jobs using "create target" will be removed.
 - "Bulk data" job option: After upgrade, jobs using "bulk data" will be turned into non-bulk data jobs.
 - Native Microsoft Excel Masking: After upgrade, MS Excel connectors, rulesets and jobs will be removed.

Licenses and Notices

The Delphix Dynamic Data Platform includes licensed, third-party products from the following companies. These products are copyrighted and all rights are reserved by the respective companies:

- Highcharts, © Highsoft

The Delphix Masking engine includes licensed, third-party products from the following companies. These products are copyrighted and all rights are reserved by the respective companies:

- Kendo UI, © Telerik

Starting with 6.0.3.0, the license info is available via a CLI/API on the engine when logged in as a system administrator.

```
engine> cd license
engine license> getLicense
engine license getLicense *> commit
```

Access to the source code of such third party open source components may be permitted or required in certain instances under the applicable open source licenses by sending an email to open-source@delphix.com.

Data Source Integration (Plugin) Release Notes

Please refer to [Data Source Integration \(Plugin\) Release Notes](#) for previous release notes.

- [Db2.x Release Notes](#)
- [EBS Release Notes](#)
- [HANA Release Notes](#)
- [PostgreSQL Release Notes](#)

Db2.x Release Notes

- [Db2 4.2.0 Release Notes](#)
- [Db2 4.1.3 Release Notes](#)
- [Db2 4.1.2 Release Notes](#)
- [Db2 4.1.1 Release Notes](#)
- [Db2 4.1.0 Release Notes](#)
- [Db2 4.0.1 Release Notes](#)
- [Db2 4.0.0 Release Notes](#)
- [Db2 3.1.1 Release Notes](#)
- [Db2 3.1.0 Release Notes](#)
- [Db2 3.0.2 Release Notes](#)
- [Db2 3.0.1 Release Notes](#)
- [Db2 3.0.0 Release Notes](#)

Db2 4.2.0 Release Notes

New Features

Dsource snapshot with User-Defined Timestamp

The Db2 Plugin version 4.2.0 introduces the dsource snapshot with the user-defined timestamp feature. This feature enables the user to specify a timestamp during dSource snapshot so that staging will roll forward to the requested timestamp before performing the dSource snapshot. For more information, see [Linking a Db2 dSource](#).

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.

• Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			

RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Plugin Version

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.2.0
--------------	---------------------------------------

3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade to 4.0.0 or 4.0.1, and follow the path below
4.0.0 4.0.1	Upgrade directly to 4.1.0, 4.1.1, 4.1.2, 4.1.3, or 4.2.0
4.1.0 4.1.1 4.1.2 4.1.3	Upgrade directly to 4.2.0

Db2 4.1.3 Release Notes

This Db2 Plugin version includes some bug fixes listed at the end of this section.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.

• Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			
RHEL 7.1			

RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Plugin Version

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.1.3
--------------	---------------------------------------

3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade to 4.0.0 or 4.0.1, and follow the path below
4.0.0 4.0.1	Upgrade directly to 4.1.0, 4.1.1, 4.1.2, or 4.1.3

Fixed Issues

Ticket Number	Summary
DB2-1812	CLI operations during snapsync may fail if commands return error messages in a different locale
DB2-1856	VDB provision fails with an issue during db2relocatedb
DB2-1875	VDB behavior when we provision them using the first snapshot vs the latter ones
DB2-1921	Fix permissions for “/toolkit_dir/DB2” along with mnts/ and logs
DB2-1904	Staging push vSDK: snapsync operation on dSource fails with TypeError: unhashable type: ‘SnapshotDefinitionDataPaths’

Db2 4.1.2 Release Notes

This Db2 Plugin version includes a bug fix listed at the end of this section.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.

• Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			
RHEL 7.1			

RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

	Plugin Version
--	----------------

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.1.2
--------------	---------------------------------------

3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade to 4.0.0 or 4.0.1, and follow the path below
4.0.0 4.0.1	Upgrade directly to 4.1.0, 4.1.1, or 4.1.2

Fixed Issues

Ticket Number	Summary
DB2-1826	Moved Staging Push redirect restore template out of the plugin and added contextual values to plugin logs
DB2-1796	Unable to get the complete content of redirect restore script in case the script size is more than 1MB
DB2-1824	DB2 VDB provisioning fails during relocateDB

Db2 4.1.1 Release Notes

The Db2 Plugin version 4.1.1 does not include any new updates.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.

• Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			
RHEL 7.1			

RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Plugin Version

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.1.1
--------------	---------------------------------------

3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade to 4.0.0 or 4.0.1, and follow the path below
4.0.0 4.0.1	Upgrade directly to 4.1.0 or 4.1.1

Db2 4.1.0 Release Notes

The Db2 Plugin version 4.1.0 includes the below improvements.

Improvements

- **User-Specified Mount Path For DB2 dSource**
Users can specify the mount path of their choice on the target host to host the dSource dataset.
- **Staging Push Automation**
A set of scripts is being provided that can be used to automate the steps related to the restore and rollforward operations on the dSource.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)



Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

	Supported DBMS Version		
Supported OS Version	10.5	11.1	11.5

RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.1.0
3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade to 4.0.0 or 4.0.1, and follow the path below
4.0.0 4.0.1	Upgrade directly to 4.1.0

Db2 4.0.1 Release Notes

The Db2 Plugin version 4.0.1 includes the below improvements.

- Added support for 11.1.4.6 fix pack.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			

RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

	Plugin Version
--	----------------

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.0.1
--------------	---------------------------------------

2.0.1	Upgrade to 2.1.0, and follow the paths below
2.1.0	Upgrade to 2.2.1, and follow the paths below
2.2.1	Upgrade to 2.3.0, and follow the paths below
2.3.0	Upgrade to 2.4.0, and follow the paths below
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below
2.5.1	Upgrade to 2.6.0, and follow the paths below
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade to 3.0.0, and follow the paths below
3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade directly to 4.0.0 or 4.0.1

Fixed Issues

Db2 plugin version 4.0.1 contains the following fixed issue.

Key	Summary
DB2-1684	[Staging Push]Enable operation on dSource should generate the Restore and Rollforward template file if it is missing

Db2 4.0.0 Release Notes

The Db2 Plugin version 4.0.0 includes the below improvements.

- Plugin generated log file will now be placed at <toolkit path>/<Delphix_COMMON>/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/logs/<instance name>.
- More user-friendly log statements will be printed.
- The logging.properties and advanceConf.config configuration files are now merged into the db2_plugin.conf file.
- Dependency on redirect restore script to get the data path information is now removed.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)



Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

	Supported DBMS Version		
Supported OS Version	10.5	11.1	11.5

RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 4.0.0
2.0.1	Upgrade to 2.1.0, and follow the paths below
2.1.0	Upgrade to 2.2.1, and follow the paths below
2.2.1	Upgrade to 2.3.0, and follow the paths below
2.3.0	Upgrade to 2.4.0, and follow the paths below
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below
2.5.1	Upgrade to 2.6.0, and follow the paths below
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade to 3.0.0, and follow the paths below
3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0, and follow the path below
3.1.0 3.1.1	Upgrade directly to the latest version

Db2 3.1.1 Release Notes

New Features

The Db2 Plugin version 3.1.1 introduces the support for full online backups, which are taken using the "exclude logs" syntax. This support is only applicable to the dSources created using the Backup and Logs ingestion mechanism on a single partition.

In Db2, users have the flexibility to exclude logs while taking the full online backups. The Db2 plugin will read the backup header from the user-provided backup file. The backup header is used to figure out whether the backup was taken using the "exclude logs" syntax or not. In case the plugin detects that it was taken using the "exclude logs" option, then the required changes will be made in the restore syntax.

Migration and Compatibility

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)



Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

	Supported DBMS Version		
Supported OS Version	10.5	11.1	11.5

RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 3.1.1
2.0.1	Upgrade to 2.1.0, and follow the paths below.
2.1.0	Upgrade to 2.2.1, and follow the paths below.
2.2.1	Upgrade to 2.3.0, and follow the paths below.
2.3.0	Upgrade to 2.4.0, and follow the paths below.
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below.
2.5.1	Upgrade to 2.6.0, and follow the paths below.
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below.
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade to 3.0.0, and follow the paths below.
3.0.0 3.0.1 3.0.2	Upgrade directly to the latest version
3.1.0	Upgrade directly to the latest version

Db2 3.1.0 Release Notes

New Features

The Db2 Plugin version 3.1.0 introduces the **Staging Push** feature. This feature will allow users to control the restore of the staging database and they can also control the state of the restored staging database.

Staging Push

With the **Staging Push** feature, the restore and rollforward operations will be managed by the user from the CLI. During dSource creation (or RESYNC) the plugin will create an **empty mount point** and will provide a restore and rollforward template that will help the user to create the redirect restore script. Once the dSource will be ready then the user can perform restore and rollforward operations from CLI. The restore must be as per the plugin provided template. After the restore, the staging database should remain in rollforward pending state, and prior to the next snapshot, the user can do a rollforward to a specific Point-in-Time (by providing a timestamp to rollforward command) against the available archive logs.

Migration and Compatibility

Supported DBMS Versions

- Db2 Enterprise Server Edition 10.5
- Db2 Advanced Enterprise Server Edition 10.5
- Db2 Enterprise Server Edition 11.1
- Db2 Advanced Enterprise Server Edition 11.1

Supported Operating Systems

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5

AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

⚠ The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗

6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 3.1.0
2.0.1	Upgrade to 2.1.0, and follow the paths below.
2.1.0	Upgrade to 2.2.1, and follow the paths below.
2.2.1	Upgrade to 2.3.0, and follow the paths below.
2.3.0	Upgrade to 2.4.0, and follow the paths below.
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below.
2.5.1	Upgrade to 2.6.0, and follow the paths below.
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below.
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade to 3.0.0, and follow the paths below.
3.0.0 3.0.1 3.0.2	Upgrade to 3.1.0

Db2 3.0.2 Release Notes

New Features

No new features. This release addresses a limitation in Db2 wherein the VDB can accept the connections when refresh, disable and delete Delphix operations are in process which eventually can corrupt the VDBs. This plugin version will reduce such occurrences by introducing maintenance window in these operations during which the VDB will not accept external connections.

Migration and Compatibility

Supported DBMS Versions

- Db2 Enterprise Server Edition 10.5
- Db2 Advanced Enterprise Server Edition 10.5
- Db2 Enterprise Server Edition 11.1
- Db2 Advanced Enterprise Server Edition 11.1

Supported Operating Systems

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)



Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
----------	---	---	---	---	---	---	---	---	---	---	---

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 3.0.2
2.0.1	Upgrade to 2.1.0, and follow the paths below.
2.1.0	Upgrade to 2.2.1, and follow the paths below.
2.2.1	Upgrade to 2.3.0, and follow the paths below.
2.3.0	Upgrade to 2.4.0, and follow the paths below.
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below.
2.5.1	Upgrade to 2.6.0, and follow the paths below.
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below.
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade directly to the latest version
3.0.0 3.0.1	Upgrade directly to the latest version

Db2 3.0.1 Release Notes

New Features

The Db2 Plugin version 3.0.1 fixes the bug wherein the jobs on UI hung intermittently due to a known IBM Db2 bug.

Migration and Compatibility

Supported DBMS Versions

- Db2 Enterprise Server Edition 10.5
- Db2 Advanced Enterprise Server Edition 10.5
- Db2 Enterprise Server Edition 11.1
- Db2 Advanced Enterprise Server Edition 11.1

Supported Operating Systems

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

-  • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
 - Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

-  Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

	Supported DBMS Version

Supported OS Version	10.5	11.1	11.5
RHEL 6.9	✓	✓	✗
RHEL 7.0	✓	✓	✗
RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Delphix Engine	Plugin Version										
	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 3.0.1
2.0.1	Upgrade to 2.1.0, and follow the paths below.
2.1.0	Upgrade to 2.2.1, and follow the paths below.
2.2.1	Upgrade to 2.3.0, and follow the paths below.
2.3.0	Upgrade to 2.4.0, and follow the paths below.
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below.
2.5.1	Upgrade to 2.6.0, and follow the paths below.
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below.
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade directly to the latest version
3.0.0	Upgrade directly to the latest version

Db2 3.0.0 Release Notes

New Features

The Db2 Plugin version 3.0.0 introduces support to Db2 Database Partitioning Feature (DPF). This will help the customer to ingest the data to a staging dsources from a Db2 DPF enabled source.

Database Partitioning Feature

Database Partitioning Feature (DPF) lets you partition your database across multiple servers. Since you can add new machines and spread your database across them, this allows users to scale their database. This means more CPUs, more memory, and more disks from each of the additional machines are available for your database. DPF can be used to manage large databases for a variety of use cases including data warehousing, data mining, online analytical processing (OLAP), or online transaction processing (OLTP) workloads.

DPF enables the user to divide a database into database partitions, a database partition is a part of a database that consists of its own data, indexes, configuration files, and transaction logs. Each database partition can be configured on the different physical server having its own set of computing resources, including CPU and storage. When a query is processed, the request is divided so each database partition processes the rows that it is responsible for. DPF can maintain consistent query performance as the table grows by providing the capability to add more processing power in the form of additional database partitions. This capability is often referred to as providing linear scalability using Db2's shared-nothing architecture.

DPF is an approach to sizing and configuring an entire database system. Please follow the recommended practices for optimal performance, reliability, and capacity growth. Please refer to IBM documentation of DPF for more details in [IBM knowledge center](#).

Pipelining logic for implementing parallel restores

The plugin employs a parallel pipeline methodology so that the restore operation of non-catalog partitions can be performed in parallel in the Database Partitioning Feature (DPF). The number of parallel restores is determined by the value of "restorePipelineLimit" (default value is 10) in <Toolkit directory>/advanceConfFlag.conf. The parameter "restorePipelineLimit" is configurable by end-users. The plugin performs parallel restore for all non-catalog partitions. For e.g. If the total number of non-catalog partitions is 15, and the "restorePipelineLimit" parameter is set to 10, the first set of 10 restores will happen in parallel. The plugin will track the restore of each partition present in the pipe. Whenever a restore of a partition completes, it will move out from the pipe and a new partition will enter into that pipe. Thus, the plugin ensures that the pipe will always have the user-configured number of partitions being restored (default=10).

Data Synchronization

When "**Customer Supplied Archive Logs**" feature is used along with Db2 DPF feature, users will have to place the archive logs inside a folder with a name as NODE<Partition number> where <Partition number> is a four-digit number. For example, if the source environment has two partitions then the user-provided log path will have folder names NODE0000 and NODE0001. Both the folders will have respective archive logs. Snapshot operation will be used to apply the archive logs.

Advance Configuration file

The advanced config file is created at <Toolkit directory>/advanceConfFlag.conf during discovery or environment refresh operation if the file does not pre-exist already. This file will have the following configurable parameters:

- **Common Group (*notCommonGroupFlag*)**
 - This parameter sets the permission of "mnts", "logs" and "code" directories at <Toolkit dir>/DB2 location.
 - Set *notCommonGroupFlag* to *false*, if the primary group of the primary environment user (user which is used to do discovery) is shared with the instance users.

- Set *notCommonGroupFlag* to *true*, if you don't want to share any group between the primary environment user and the instance users.
- Once the necessary changes are made, refresh the environment in order to make the changes applicable.

By default, the *notCommonGroupFlag* parameter is not specified in the file. This means that the plugin implicitly assumes its value as false. The valid values for this parameter are true or false.

- **Allow Source Database on Same Instance (*allowSourceDBonSameInst*)**

By default, the DB2 Plugin restricts the provisioning of a VDB on a DB2 instance which contains a database with the same name as the VDB's source database

For example:

- When provisioning a VDB from a VDB then source VDB is treated as the source database.
- When provisioning a VDB from a dSource then a staging database is treated as the source database.

Set *allowSourceDBonSameInst* to true, if the user wants to provision a VDB on an instance which contains a database with the same name as the VDB's source database. The valid values for this parameter are true or false.

- **Restore Pipeline limit (*restorePipelineLimit*)**

Db2 DPF Plugin performs restoration of all non-catalog partitions in parallel.

The parameter *restorePipelineLimit* allows the user to configure the number of partitions that should be restored in parallel. When the new *advanceConfFlag.conf* file is created during environment refresh or discovery operation, the default value of *restorePipelineLimit* is set to 10. If the *advanceConfFlag.conf* file is already present, the user will have to set the value of *restorePipelineLimit* parameter to the desired value. Valid value for this field is any positive integer.

A default config file created during discovery or environment refresh operation:

```
# If the user is not sharing a common group between primary environment user and
instance user.

# Then the user needs to set parameter notCommonGroupFlag as true.

# notCommonGroupFlag=true

# During provision operation check if instance contains a database name which is
identical to the source DB name used for provisioning operation.

# By default plugin will never allow creating a VDB on the instance where source
database (or other database which is identical in name with source DB) already exists.

# If the user still wants to create a VDB on the instance which contains a database
name which is identical to the source DB name used for provisioning operation,

# then the user needs to set parameter allowSourceDBonSameInst as true.

# allowSourceDBonSameInst=true

# Limit for parallel restores

restorePipelineLimit=10
```

Migration and Compatibility

Supported DBMS Versions

- Db2 Enterprise Server Edition 10.5
- Db2 Advanced Enterprise Server Edition 10.5
- Db2 Enterprise Server Edition 11.1
- Db2 Advanced Enterprise Server Edition 11.1

Supported Operating Systems

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.

• Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			

RHEL 7.1	✓	✓	✗
RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

	Plugin Version
--	----------------

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path

Use the table below to determine the most efficient upgrade path from your current version to the latest version of the Db2 Plugin.

Your Version	Recommended upgrade path to Db2 3.0.0
--------------	---------------------------------------

2.0.1	Upgrade to 2.1.0, and follow the paths below.
2.1.0	Upgrade to 2.2.1, and follow the paths below.
2.2.1	Upgrade to 2.3.0, and follow the paths below.
2.3.0	Upgrade to 2.4.0, and follow the paths below.
2.4.0 2.4.1 2.4.2	Upgrade to 2.5.1, and follow the paths below.
2.5.1	Upgrade to 2.6.0, and follow the paths below.
2.6.0 2.6.1	Upgrade to 2.7.0, and follow the paths below.
2.7.0 2.7.1 2.7.2 2.7.3	Upgrade to 3.0.0

Known Issues

Db2 plugin version 3.0.0 contains the following known issue:

Key	Summary	Workaround / Comments
DB2-1318	Jobs on UI might hang intermittently due to a known IBM Db2 bug	This issue will be only observed when multiple datasets jobs, within the same Db2 instance, will get executed at the same time.

EBS Release Notes

- [EBS 12.2 Plugin Release Notes](#)

EBS 12.2 Plugin Release Notes

- [EBS 12.2 Plugin 4.0.0 Release Notes](#)
- [EBS 12.2 Plugin 3.0.3 Release Notes](#)
- [EBS 12.2 Plugin 3.0.2 Release Notes](#)
- [EBS 12.2 Plugin 3.0.1 Release Notes](#)
- [EBS 12.2 Plugin 3.0.0 Release Notes](#)
- [EBS 12.2 Plugin 2.0.1 Release Notes](#)
- [EBS 12.2 Plugin 2.0.0 Release Notes](#)
- [EBS 12.2 Plugin 1.8.4 Release Notes](#)
- [EBS 12.2 Plugin 1.8.3 Release Notes](#)
- [EBS 12.2 Plugin 1.8.2 Release Notes](#)

EBS 12.2 Plugin 4.0.0 Release Notes

What's New?

The EBS 12.2 plugin version 4.0.0 release is based on the python-based SDK feature and has the following enhancements:

- The plugin-generated logs are now available on Delphix Engine and will be part of the support bundle. The logs for provisioning operations will also be available on the remote host. You will get the flexibility to set the log level, log file size, and log retentions for remote host side logging through the `ebs_plugin.conf` file.
- The plugin discovers the application on the source appsTier environment during the discovery operation.
- New password fields are available to the Delphix Engine users.
- This release adds support for EBS 12.2 with 12C on Solaris SPARC.
- Context file-based `adcfgclone` for EBS 12.2 with 19C database as per the Oracle Documentation.
- This release improves hooks for database provisioning.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)
- Oracle19c MT (Minimum DB version supported 19.3.0.0)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		
RHEL 6.7		
RHEL 6.8		
RHEL 6.9		
RHEL 7.0		
RHEL 7.1		

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 7.2	✓	N/A
RHEL 7.3	✓	N/A
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

Reference

Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗

	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c											
Supported Delphix Engine	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0		
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
1.8.3	
1.8.4	
2.0.0	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.
2.0.1	

3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Related Topics

[EBS 12.2 Plugin Release Notes](#)

EBS 12.2 Plugin 3.0.3 Release Notes

What's New?

This is a patch release that aims to provide the bug fixes that are listed in the Fixed Issues section at the end of this page.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)
- Oracle19c MT (Minimum DB version supported 19.3.0.0)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes
	No

Color	Supported?
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		
RHEL 6.7		
RHEL 6.8		
RHEL 6.9		
RHEL 7.0		
RHEL 7.1		
RHEL 7.2		
RHEL 7.3		

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

i Reference

Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗

	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c											
Supported Delphix Engine	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0		
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2 1.8.3 1.8.4	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
2.0.0 2.0.1	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.

3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Fixed Issues

Key	Summary
EBS-1027	EBS 12.2 Plugin will validate user-provided ORACLE_HOME during DbTier provisioning
EBS-1028	Enhancements identified in dbTechStack code
EBS-1038	Fix shellchecks warnings/errors in hooks and sudoers entry for target dbtier
EBS-1066	Make the check stricter for listener process in DBTechStack status module and fix shellchecks

EBS 12.2 Plugin 3.0.2 Release Notes

What's New?

This is a patch release that aims to provide the following updates to EBS 12.2 plugin.

- Updated Hooks for 19c MT low and high privileged users.
- Fixed an issue where 19c DBTechStack provision fails during `adcfgclone` with java exception.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)
- Oracle19c MT (Minimum DB version supported 19.3.0.0)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes

Color	Supported?
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		
RHEL 6.7		
RHEL 6.8		
RHEL 6.9		
RHEL 7.0		
RHEL 7.1		
RHEL 7.2		

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 7.3	✓	N/A
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

Reference

Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗

	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c											
Supported Delphix Engine	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0		
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
1.8.3	
1.8.4	
2.0.0	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.
2.0.1	

3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Fixed Issues

Key	Summary
EBS-890	19C DBTechStack provision fails during <code>adcfgclone</code> with java exception

EBS 12.2 Plugin 3.0.1 Release Notes

What's New?

This is a patch release that aims to provide the following updates to EBS 12.2 plugin.

- Added support for RHEL 8.3
- Updated Hooks for 19c MT low and high privileged users.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)
- Oracle19c MT (Minimum DB version supported 19.3.0.0)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes

Color	Supported?
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		
RHEL 6.7		
RHEL 6.8		
RHEL 6.9		
RHEL 7.0		
RHEL 7.1		
RHEL 7.2		

	Supported EBS 12.2 Apps version	
Supported OS Version	Oracle 11g/12c	Oracle 19c
RHEL 7.3	✓	N/A
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

Reference

Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗

	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c											
Supported Delphix Engine	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0		
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
1.8.3	
1.8.4	
2.0.0	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.
2.0.1	

3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Known Issues

Key	Summary
EBS-836	"ebs_<PDB_SID>" listener services are getting disappeared on VDB stop/start. As a workaround, a new Third Configure Clone Hook has been added.

EBS 12.2 Plugin 3.0.0 Release Notes

New Features

EBS Plugin version 3.0.0 aims to support Oracle 19c MT architecture. Moreover, a new field has been added on the provisioning UI. This will provide the flexibility to the end-user to provide the maximum time allowed to virtualize the application tier.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)
- Oracle19c MT (Minimum DB version supported 19.3.0.0)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_3.0.0
5.3.0.0	No
5.3.1.0	No
5.3.2.0	No

5.3.3.0	No
5.3.4.0	No
5.3.5.0	No
5.3.6.0	No
5.3.7.0	No
5.3.8.0-5.3.8.1	No
5.3.9.0	No
6.0.0.0	No
6.0.1.0	No
6.0.2.0	No
6.0.3.0	No
6.0.4.0	No
6.0.5.0	No
6.0.6.0	No
6.0.7.0	No
6.0.8.0	No
6.0.9.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 3.0.0 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

Your Version	Recommended Upgrade Path to 3.0.0
1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, 1.8.2, 1.8.3, or 1.8.4 and follow the upgrade path below.
1.8.0 1.8.1 1.8.2 1.8.3 1.8.4	Upgrade to 2.0.0 or 2.0.1, and follow the upgrade path below.
2.0.0 2.0.1	Upgrade to 3.0.0, the latest plugin version.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 3.0.0
 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0.

- Customers should skip a minor Plugin 1.3.1 and 1.8.1/1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0 .

Key points:

- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- SSL enabled Environment

- V2P

Known Issues

Key	Summary
EBS-836	"ebs_<PDB_SID>" listener services are getting disappeared on VDB stop/start. As a workaround, a new configure clone hook has been added.

Fixed Issues

EBS plugin version 3.0.0 contains the following bug fixes.

Key	Summary
EBS-767	EBS appsTier linking failed while identifying Delta patch level information from low-privilege OS user
EBS-779	Continued timeouts - toolkit.ebs122-app-1-8-2.errorTimeout {admanagedsrvctl}
EBS-787	Toolkit should perform prior dependency/pre-requisite check in target host before starting to provision dBtechStack and Appstier
EBS-788	Error not being shown to the customer
EBS-798	GUI shows vfiles ORACLE_HOME as stopped however it is not
EBS-778	Expect timeout set to 8 hours which is too short
EBS-798	GUI shows vfiles ORACLE_HOME as stopped however it is not
EBS-816	EBS 12.2 disable Appstier vdb failed with Delphix cannot stop all processes under /u01/oracle/VIS on host

EBS 12.2 Plugin 2.0.1 Release Notes

New Features

EBS Plugin version 2.0.1 aims to successfully stop application services while provisioning or refresh operation. Hence, creating successful snapshots.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11
- Advanced Interactive eXecutive (AIX) – See the matrix at the bottom of the [Oracle E-Business Suite Support Matrix](#).

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_2.0.1
5.3.0.0	No
5.3.1.0	Yes
5.3.2.0	Yes

5.3.3.0	Yes
5.3.4.0	Yes
5.3.5.0	Yes
5.3.6.0	Yes
5.3.7.0	Yes
5.3.8.0-5.3.8.1	Yes
5.3.9.0	Yes
6.0.0.0	Yes
6.0.1.0	Yes
6.0.2.0	Yes
6.0.3.0	Yes
6.0.4.0	Yes
6.0.5.0	Yes
6.0.6.0	Yes
6.0.7.0	Yes
6.0.8.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 2.0.1 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

<i>Your Version</i>	<i>Recommended Upgrade Path to 2.0.1</i>
----------------------------	---

1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, 1.8.2, 1.8.3, or 1.8.4 and follow the upgrade path below.
1.8.0 to 1.8.4	Upgrade to 2.0.0 or 2.0.1, the latest plugin version.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 2.0.0 or 2.0.1 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0 or 2.0.1.

- Customers should skip a minor Plugin 1.3.1 and 1.8.1/1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0 or 2.0.1.

Key points:

- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- SSL enabled Environment
- V2P

Known Issues

N/A

Fixed Issues

EBS plugin version 2.0.1 contains the following features and bug fixes:

Key	Summary
-----	---------

EBS-783	Provisioning VDB from Appstier dSource fails with adpreclone connection with Oracle WebLogic Administration Server
---------	--

EBS 12.2 Plugin 2.0.0 Release Notes

New Features

EBS Plugin version 2.0.0 aims to implement Sudo Privileges feature for both EBS 12.2 Database Tier and Application Tier.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11
- Advanced Interactive eXecutive (AIX) – See the matrix at the bottom of the [Oracle E-Business Suite Support Matrix](#).

Plugin/Delphix Engine Compatibility

The Plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_2.0.0
5.3.0.0	No
5.3.1.0	Yes
5.3.2.0	Yes

5.3.3.0	Yes
5.3.4.0	Yes
5.3.5.0	Yes
5.3.6.0	Yes
5.3.7.0	Yes
5.3.8.0-5.3.8.1	Yes
5.3.9.0	Yes
6.0.0.0	Yes
6.0.1.0	Yes
6.0.2.0	Yes
6.0.3.0	Yes
6.0.4.0	Yes
6.0.5.0	Yes
6.0.6.0	Yes
6.0.7.0	Yes
6.0.8.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 2.0.0 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

<i>Your Version</i>	<i>Recommended Upgrade Path to 2.0.0</i>
----------------------------	---

1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, 1.8.2, 1.8.3, or 1.8.4 and follow the upgrade path below.
1.8.0 to 1.8.4	Upgrade to 2.0.0, the latest plugin version.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 2.0.0
 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0.

- Customers should skip a minor Plugin 1.3.1 and 1.8.1/1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4 → 2.0.0.

Key points:

- Plugin 2.0.0 supports IDEA-1072: Delphix OS account ("delphix_os") on source/target servers should obtain privileges only through "sudo".
- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- SSL enabled Environment
- V2P

Known Issues

N/A

Fixed Issues

EBS plugin version 1.8.4 contains the following features and bug fixes:

Key	Summary
-----	---------

EBS-757	EBS plugin should check for data accessibility during status check
---------	--

EBS 12.2 Plugin 1.8.4 Release Notes

New Features

EBS Plugin version 1.8.4 aims to implement Weblogic password complexity validation as per the Oracle standards.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)

 Oracle does not yet support the multi-tenant option for E-Business Suites, so Delphix cannot support this either.

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11
- Advanced Interactive eXecutive (AIX) – See the matrix at the bottom of the [Oracle E-Business Suite Support Matrix](#).

Plugin/Delphix Engine Compatibility

The Plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_1.8.4
5.3.0.0	No
5.3.1.0	Yes

5.3.2.0	Yes
5.3.3.0	Yes
5.3.4.0	Yes
5.3.5.0	Yes
5.3.6.0	Yes
5.3.7.0	Yes
5.3.8.0-5.3.8.1	Yes
5.3.9.0	Yes
6.0.0.0	Yes
6.0.1.0	Yes
6.0.2.0	Yes
6.0.3.0	Yes
6.0.4.0	Yes
6.0.5.0	Yes
6.0.6.0	Yes
6.0.7.0	Yes
6.0.8.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 1.8.4 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

Your Version	Recommended Upgrade Path to 1.8.3
1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, 1.8.2, 1.8.3, or 1.8.4 and follow the upgrade path below.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 1.8.x
 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4

- Customers should skip a minor Plugin 1.3.1 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, 1.8.3. or 1.8.4

Key points:

- Plugin 1.8.4 is a patch release
- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- Sudo Privilege Plugin for EBS 12.2 release
- Support SSL enabled Environment
- V2P

Known Issues

N/A

Fixed Issues

EBS plugin version 1.8.4 contains the following features and bug fixes:

Key	Summary
EBS-712	WLS Password Complexity Issue - Merge ESCL_2983 into EBS 12.2 Plugin.
EBS-713	Multi-line support for SQL query O/P - merge ESCL-2988 into EBS 12.2 Plugin.

EBS 12.2 Plugin 1.8.3 Release Notes

New Features

EBS Plugin version 1.8.3 introduces a new feature while refreshing EBS appsTier where a Pre Refresh hook will help in aborting concurrent manager services and reduce the time taken to gracefully stop the EBS application.. The password required to execute the same will be pulled from a text file present on the target appsTier host which will be created automatically during appsTier provisioning.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)

 Oracle does not yet support the multi-tenant option for E-Business Suites, so Delphix cannot support this either.

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11
- Advanced Interactive eXecutive (AIX) – See the matrix at the bottom of the [Oracle E-Business Suite Support Matrix](#).

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_1.8.3
5.3.0.0	No

5.3.1.0	Yes
5.3.2.0	Yes
5.3.3.0	Yes
5.3.4.0	Yes
5.3.5.0	Yes
5.3.6.0	Yes
5.3.7.0	Yes
5.3.8.0-5.3.8.1	Yes
5.3.9.0	Yes
6.0.0.0	Yes
6.0.1.0	Yes
6.0.2.0	Yes
6.0.3.0	Yes
6.0.4.0	Yes
6.0.5.0	Yes
6.0.6.0	Yes
6.0.7.0	Yes
6.0.8.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 1.8.3 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

<i>Your Version</i>	<i>Recommended Upgrade Path to 1.8.3</i>
1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, 1.8.2, or 1.8.3 and follow the upgrade path below.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 1.8.x
 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, or 1.8.3.

- Customers should skip a minor Plugin 1.3.1 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1, 1.8.2, or 1.8.3.

Key points:

- Plugin 1.8.3 is a patch release
- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- Sudo Privilege Plugin for EBS 12.2 release
- Support SSL enabled Environment
- V2P
- BYOOJ

Known Issues

EBS plugin version 1.8.3 contains the following known issues :

Key	Summary
-----	---------

EBS-692	EBS link operation fails when source EBS environment is linked with BYOOJ
DLPX-69617	OperationRunnerImpl incorrectly assumes toolkit JDK always exists

Fixed Issues

EBS plugin version 1.8.3 contains the following features and bug fixes:

Key	Summary
EBS-663	Abort EBS ConcMgr services before stopping EBS appsTier services during refresh.

EBS 12.2 Plugin 1.8.2 Release Notes

New Features

EBS Plugin version 1.8.2 introduces a new feature where the user will not need to specify the source apps schema password in EBS DB hooks provision wizard. This password can now be pulled from a text file present on the target DB host which will be created automatically during DBTechStack provisioning.

Plugin Installation

Plugins to support EBS virtualization are not automatically included in the Virtualization Engine deployment and can be installed by the end-user after uploading the plugin through the Manage > Plugin screen.

Migration and Compatibility

Supported Versions

Oracle SI dbTechStack and Database

EBS 12.2

- Oracle11gR2 (Minimum DB version supported 11.2.0.4)
- Oracle12cR1 (Minimum DB version supported 12.1.0.2)

 Oracle does not yet support the multi-tenant option for E-Business Suites, so Delphix cannot support this either.

Oracle EBS appsTier

- Single-node appsTier
- Multi-node appsTier with shared APPL_TOP

 The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

Supported Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris – SPARC 9, 10, 11
- Advanced Interactive eXecutive (AIX) – See the matrix at the bottom of [Oracle E-Business Suite Support Matrix](#).

Plugin/Delphix Engine Compatibility

The Plugin should be installed on compatible Delphix Engines per the table below. For plugin/engine compatibility of all the released plugin versions, see [EBS Matrix](#).

For EBS 12.2 customers:

Delphix Engine	EBS_1.8.2
5.3.0.0	No

5.3.1.0	Yes
5.3.2.0	Yes
5.3.3.0	Yes
5.3.4.0	Yes
5.3.5.0	Yes
5.3.6.0	Yes
5.3.7.0	Yes
5.3.8.0-5.3.8.1	Yes
5.3.9.0	Yes
6.0.0.0	Yes
6.0.1.0	Yes
6.0.2.0	Yes
6.0.3.0	Yes
6.0.4.0	Yes
6.0.5.0	Yes
6.0.6.0	Yes
6.0.7.0	Yes
6.0.8.0	Yes

Plugin Upgrade Path

Path A - New plugin installation for EBS customers:

- Upload 1.8.2 plugin for EBS 12.2 customers

Path B - Existing EBS customers with different plugin version:

For EBS 12.2 customers:

<i>Your Version</i>	<i>Recommended Upgrade Path to 1.8.2</i>
1.0.0	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0 1.3.1	Upgrade to 1.4.0 and follow the upgrade path below.
1.4.0	Upgrade to 1.5.0 and follow the upgrade path below.
1.5.0	Upgrade to 1.6.0 and follow the upgrade path below.
1.6.0	Upgrade to 1.7.0 and follow the upgrade path below.
1.7.0	Upgrade to 1.8.0, 1.8.1, or 1.8.2 and follow the upgrade path below.

Example : Upgrade path for EBS 12.2 customer with existing Plugin version 1.1.0 who wants to upgrade to 1.8.1
 Plugin version : 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1 or 1.8.2

- Customers should skip a minor Plugin 1.3.1 version upgrade.
- If EBS 12.2 customers are on Plugin version 1.3.1, the plugin upgrade path will be : 1.3.1 → 1.4.0 → 1.5.0 → 1.6.0 → 1.7.0 → 1.8.0 or 1.8.1 or 1.8.2

Key points:

- Plugin 1.8.2 is a patch release that will support EBS-597 Delphix Hooks Password-Storing Apps password to file for EBS appsTier.
- Customers can skip a patch plugin release upgrade if they intend to upgrade their plugin to a higher version. The Delphix Engine simply ignores patch versions when handling plugin upgrade logic. It only validates the major/minor plugin version during the plugin upgrade.

Unsupported EBS Versions and Features

- Sudo Privilege Plugin for EBS 12.2 release
- Support SSL enabled Environment
- V2P
- BYOOJ

Fixed Issues

EBS plugin version 1.8.2 contains the following features and bug fixes:

Key	Summary
-----	---------

EBS_566	Display port pool number in the error toolkit.ebs122-app-1-7-0.portPoolNotFreeError.
EBS-597	Delphix Hooks Password-Storing Apps password to file for EBS appsTier

Known Issues

Key	Summary
EBS-692	EBS link operation fails when source EBS environment is linked with BYOOJ
DLPX-69617	OperationRunnerImpl incorrectly assumes toolkit JDK always exists

HANA Release Notes

- [HANA 2.0 Plugin Version 6.2.0 Release Notes](#)
- [HANA 2.0 Plugin Version 6.1.1 Release Notes](#)
- [HANA 2.0 Plugin Version 6.1.0 Release Notes](#)
- [HANA 2.0 Plugin Version 6.0.1 Release Notes](#)
- [HANA 2.0 Plugin Version 6.0.0 Release Notes](#)
- [HANA 2.0 Plugin Version 4.3.2 Release Notes](#)
- [HANA 2.0 Plugin Version 4.3.1 Release Notes](#)
- [HANA 2.0 Plugin Version 4.3.0 Release Notes](#)
- [HANA 2.0 Plugin Version 4.2.0 Release Notes](#)
- [HANA 2.0 Plugin Version 4.1.2 Release Notes](#)
- [HANA 2.0 Plugin Version 4.1.1 Release Notes](#)

HANA 2.0 Plugin Version 6.2.0 Release Notes

New Features

The HANA 2.0 Plugin version 6.2.0 introduces the **Staging Push** feature. This feature provides a new data ingestion mechanism that helps users to push data into the Delphix provided mount point on their own.

Plugin Installation

Plugins to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

This section describes the HANA (DBMS) versions that are supported by the HANA 2.0 Plugin, as well as the compatible operating systems (OS), for use on the source and target environments.

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Supported DBMS Versions

HANA Version (Standard and Enterprise Editions)	Support Package Stack (SPS)
2.0	04
2.0	05

 **Important:** A target may have a higher SPS release than a source; the reverse is not supported.

Key:

Color	Supported?
	Yes
	No

Supported Operating System

Supported DBMS Version

Supported OS Version	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
RHEL 8.1	✓	✓
RHEL 7.3/7.4/7.5	✓	✗
CentOS 7.3/7.4/7.5	✓	✗
RHEL 7.6	✓	✓
CentOS 7.6	✓	✓
SUSE 12 SP 03	✓	✗
SUSE 15 SP 02	✗	✓

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	Lua-based plugin						vSDK-based Plugin						
	4.1.1	4.1.2	4.2.0	4.3.0	4.3.1	4.3.2	6.0.0	6.0.1	6.1.0	6.1.1	6.2.0	6.3.0	
6.0.0.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.1.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.2.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.3.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.4.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.5.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	

	Lua-based plugin						vSDK-based Plugin					
6.0.6.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point-in-Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Although the Staging Push solution is generic and should work with HANA Streaming Replication (HSR) but for this release, we have not validated the same.
- Third-party backups tools are not supported with Pull based ingestion mechanism.
- More than one instance of the same HANA service is not supported with this release.
- Switching between the Staging Push to Pull model is not supported with this release.

Plugin Upgrade Path

Upgrade from HANA Direct (Lua-based) plugin to HANA Staged (VSDK-based) plugin is not supported.

To upgrade from the existing HANA Staged (VSDK-based) plugin follow the recommended path below.

Your Version	Recommended upgrade path to 6.2.0
6.0.0 6.0.1	Upgrade to 6.1.0 and follow the path below

Your Version	Recommended upgrade path to 6.2.0
6.1.0 6.1.1	Upgrade directly to 6.2.0

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
DLPX-75383	Refresh/Disable not working after failed delete operation	Force delete the VDB. Manually remove the mount point
DLPX-75385	Undo Refresh operation after a failed Refresh job leads to 502 Bad gateway and management stack restart	None
DLPX-67911	Failed enable operation leaving a mount point and next delete operation creating a stale file handle	Manually remove the mount point
CTCHANA-1599	Mount Point change not detected by the plugin may lead to multiple mounts	Mount point not in use can be removed manually
CTCHANA-1598	Plugin loses control in case of user cancel operation.	Manual cleanup can be done in case the plugin leaves any tenant database/ files on the VM
CTCHANA-1586	Refresh after a Force disable on a VDB fails with 'mount not empty' error	Manually remove the mount before proceeding with the Refresh
CTCHANA-1565	Stop/Start operation on a VDB post failed post_refresh hook fails with inappropriate error	None
CTCHANA-1514	dSource tenant database can be accessed and written onto	None
CTCHANA-1610	'test_force_disable_enable_vdb' fails across multiple images	None

HANA 2.0 Plugin Version 6.1.1 Release Notes

This patch release contains an enhancement that allows the `create_backup_metadata.sh` script to be executed in a non-interactive mode as well.

Plugin Installation

Plugins to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

This section describes the HANA (DBMS) versions that are supported by the HANA 2.0 Plugin, as well as the compatible operating systems (OS), for use on the source and target environments.

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Supported DBMS Versions

HANA Version (Standard and Enterprise Editions)	Support Package Stack (SPS)
2.0	04
2.0	05

 **Important:** A target may have a higher SPS release than a source; the reverse is not supported.

Key:

Color	Supported?
	Yes
	No

Supported Operating System

	Supported DBMS Version	
Supported OS Version	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)

RHEL 8.1	✓	✓
RHEL 7.3/7.4/7.5	✓	✗
CentOS 7.3/7.4/7.5	✓	✗
RHEL 7.6	✓	✓
CentOS 7.6	✓	✓
SUSE 12 SP 03	✓	✗
SUSE 15 SP 02	✗	✓

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	Lua-based plugin						vSDK-based Plugin						
	4.1.1	4.1.2	4.2.0	4.3.0	4.3.1	4.3.2	6.0.0	6.0.1	6.1.0	6.1.1	6.2.0	6.3.0	
6.0.0.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.1.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.2.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.3.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.4.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.5.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.6.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.7.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	

	Lua-based plugin						vSDK-based Plugin					
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point-in-Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Although the Staging Push solution is generic and should work with HANA Streaming Replication (HSR) but for this release, we have not validated the same.
- Third-party backups tools are not supported with Pull based ingestion mechanism.
- More than one instance of the same HANA service is not supported with this release.
- Switching between the Staging Push to Pull model is not supported with this release.

Plugin Upgrade Path

Upgrade from HANA Direct (Lua-based) plugin to HANA Staged (VSDK-based) plugin is not supported.

To upgrade from the existing HANA Staged (VSDK-based) plugin follow the recommended path below.

Your Version	Recommended upgrade path to 6.1.1
6.0.0 6.0.1	Upgrade to 6.1.0 and follow the path below
6.1.0	Upgrade directly to 6.1.0 or 6.1.1

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
DLPX-75383	Refresh/Disable not working after failed delete operation	Force delete the VDB. Manually remove the mount point
DLPX-75385	Undo Refresh operation after a failed Refresh job leads to 502 Bad gateway and management stack restart	None
DLPX-67911	Failed enable operation leaving a mount point and next delete operation creating a stale file handle	Manually remove the mount point
CTCHANA-1599	Mount Point change not detected by the plugin may lead to multiple mounts	Mount point not in use can be removed manually
CTCHANA-1598	Plugin loses control in case of user cancel operation.	Manual cleanup can be done in case the plugin leaves any tenant database/ files on the VM
CTCHANA-1586	Refresh after a Force disable on a VDB fails with 'mount not empty' error	Manually remove the mount before proceeding with the Refresh
CTCHANA-1565	Stop/Start operation on a VDB post failed post_refresh hook fails with inappropriate error	None
CTCHANA-1514	dSource tenant database can be accessed and written onto	None
CTCHANA-1610	'test_force_disable_enable_vdb' fails across multiple images	None

Fixed Issues

The following issues are fixed in this release:

Delphix Ticket Number	Summary
CTCHANA-1661	v6.0.0.0 create_backup_metadata.sh needs ability to run non-interactively
CTCHANA-1667	HANA Recovery failure on staging host yields difficult to decipher error

HANA 2.0 Plugin Version 6.1.0 Release Notes

Features

HANA Port Control

With this release, a solution has been introduced that keeps the port numbers consistent throughout the VDB life cycle so that the connections made to the VDBs are not disrupted during their life cycle.

Plugin Installation

Plugins to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

This section describes the HANA (DBMS) versions that are supported by the HANA 2.0 Plugin, as well as the compatible operating systems (OS), for use on the source and target environments.

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Supported DBMS Versions

HANA Version (Standard and Enterprise Editions)	Support Package Stack (SPS)
2.0	04
2.0	05

 **Important:** A target may have a higher SPS release than a source; the reverse is not supported.

Key:

Color	Supported?
	Yes
	No

Supported Operating System

Supported DBMS Version

Supported OS Version	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
RHEL 8.1	✓	✓
RHEL 7.3/7.4/7.5	✓	✗
CentOS 7.3/7.4/7.5	✓	✗
RHEL 7.6	✓	✓
CentOS 7.6	✓	✓
SUSE 12 SP 03	✓	✗
SUSE 15 SP 02	✗	✓

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	Lua-based plugin						vSDK-based Plugin						
	4.1.1	4.1.2	4.2.0	4.3.0	4.3.1	4.3.2	6.0.0	6.0.1	6.1.0	6.1.1	6.2.0	6.3.0	
6.0.0.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.1.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.2.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.3.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.4.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.5.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	

	Lua-based plugin						vSDK-based Plugin					
6.0.6.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point-in-Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Although the Staging Push solution is generic and should work with HANA Streaming Replication (HSR) but for this release, we have not validated the same.
- Third-party backups tools are not supported with Pull based ingestion mechanism.
- More than one instance of the same HANA service is not supported with this release.
- Switching between the Staging Push to Pull model is not supported with this release.

Plugin Upgrade Path

Upgrade from HANA Direct (Lua-based) plugin to HANA Staged (VSDK-based) plugin is not supported.

To upgrade from the existing HANA Staged (VSDK-based) plugin follow the recommended path below.

Your Version	Recommended upgrade path to 6.1.0
6.0.0 6.0.1	Upgrade to 6.1.0 and follow the path below

Your Version	Recommended upgrade path to 6.1.0
6.1.0	Upgrade directly to 6.1.0

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
DLPX-75383	Refresh/Disable not working after failed delete operation	Force delete the VDB. Manually remove the mount point
DLPX-75385	Undo Refresh operation after a failed Refresh job leads to 502 Bad gateway and management stack restart	None
DLPX-67911	Failed enable operation leaving a mount point and next delete operation creating a stale file handle	Manually remove the mount point
CTCHANA-1599	Mount Point change not detected by the plugin may lead to multiple mounts	Mount point not in use can be removed manually
CTCHANA-1598	Plugin loses control in case of user cancel operation.	Manual cleanup can be done in case the plugin leaves any tenant database/ files on the VM
CTCHANA-1586	Refresh after a Force disable on a VDB fails with 'mount not empty' error	Manually remove the mount before proceeding with the Refresh
CTCHANA-1565	Stop/Start operation on a VDB post failed post_refresh hook fails with inappropriate error	None
CTCHANA-1514	dSource tenant database can be accessed and written onto	None
CTCHANA-1610	'test_force_disable_enable_vdb' fails across multiple images	None

Fixed Issues

The following issue is fixed in this release:

Delphix Ticket Number	Summary
CTCHANA-821	HANA SQL Port value should remain consistent for the life of VDB.

HANA 2.0 Plugin Version 6.0.1 Release Notes

This is a bug-fix release of the HANA plugin version 6.0.1.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS04
- HANA Standard and Enterprise Edition 2.0 SPS05

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- SUSE Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version		HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
	Supported OS Version		
	RHEL 7.3	Supported	Not Supported
	RHEL 7.4	Supported	Not Supported
	RHEL 7.5	Supported	Not Supported
	RHEL 7.6	Supported	Supported
	RHEL 8.1	Supported	Supported

SUSE Linux

Supported DBMS Version	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)

Supported OS Version		
SUSE 12 SP3	Supported	Not Supported
SUSE 15 SP2	Not Supported	Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 6.0.0	HANA 6.0.1
6.0.0.x	No	No
6.0.1.x	No	No
6.0.2.x	No	No
6.0.3.x	No	No
6.0.4.x	No	No
6.0.5.x	No	No
6.0.6.x	No	No
6.0.7.x	No	No
6.0.8.x	Yes	Yes
6.0.9.x	Yes	Yes

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Upgrade from HANA Direct (Lua-based) plugin to HANA Staged (VSDK-based) plugin is not supported.

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
DLPX-75383	Refresh/Disable not working after failed delete operation	Force delete the VDB. Manually remove the mount point
DLPX-75385	Undo Refresh operation after a failed Refresh job leads to 502 Bad gateway and management stack restart	None
DLPX-67911	Failed enable operation leaving a mount point and next delete operation creating a stale file handle	Manually remove the mount point
CTCHANA-1599	Mount Point change not detected by the plugin may lead to multiple mounts	Mount point not in use can be removed manually
CTCHANA-1598	Plugin loses control in case of user cancel operation.	Manual cleanup can be done in case the plugin leaves any tenant database/ files on the VM
CTCHANA-1586	Refresh after a Force disable on a VDB fails with 'mount not empty' error	Manually remove the mount before proceeding with the Refresh
CTCHANA-1565	Stop/Start operation on a VDB post failed post_refresh hook fails with inappropriate error	None
CTCHANA-1514	Dsource tenant database can be accessed and written onto	None
CTCHANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None
CTCHANA-1610	'test_force_disable_enable_vdb' fails across multiple images	None

HANA 2.0 Plugin Version 6.0.0 Release Notes

This is the release of the HANA plugin version 6.0.0.

New Features

This is a conversion from HANA direct plugin to vSDK-based, staged plugin. This plugin will support only customer-provided/External backups ingestion mechanism.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS04
- HANA Standard and Enterprise Edition 2.0 SPS05

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- SUSE Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version		HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version			
	RHEL 7.3	Supported	Not Supported
	RHEL 7.4	Supported	Not Supported
	RHEL 7.5	Supported	Not Supported
	RHEL 7.6	Supported	Supported
	RHEL 8.1	Supported	Supported

SUSE Linux

Supported DBMS Version		HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version			
SUSE 12 SP3	Supported	Not Supported	
SUSE 15 SP2	Not Supported	Supported	

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 6.0.0
6.0.0.x	No
6.0.1.x	No
6.0.2.x	No
6.0.3.x	No
6.0.4.x	No
6.0.5.x	No
6.0.6.x	No
6.0.7.x	No
6.0.8.x	Yes
6.0.9.x	Yes

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.

- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Upgrade from HANA Direct (Lua-based) plugin to HANA Staged (VSDK-based) plugin is not supported.

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
DLPX-75383	Refresh/Disable not working after failed delete operation	Force delete the VDB. Manually remove the mount point
DLPX-75385	Undo Refresh operation after a failed Refresh job leads to 502 Bad gateway and management stack restart	None
DLPX-67911	Failed enable operation leaving a mount point and next delete operation creating a stale file handle	Manually remove the mount point
CTCHANA-1599	Mount Point change not detected by the plugin may lead to multiple mounts	Mount point not in use can be removed manually
CTCHANA-1598	Plugin loses control in case of user cancel operation.	Manual cleanup can be done in case the plugin leaves any tenant database/ files on the VM
CTCHANA-1586	Refresh after a Force disable on a VDB fails with 'mount not empty' error	Manually remove the mount before proceeding with the Refresh
CTCHANA-1565	Stop/Start operation on a VDB post failed post_refresh hook fails with inappropriate error	None
CTCHANA-1514	Dsource tenant database can be accessed and written onto	None
CTCHANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None

HANA 2.0 Plugin Version 4.3.2 Release Notes

This is a patch release of the HANA plugin version 4.3.2.

New Features

No new features.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS04
- HANA Standard and Enterprise Edition 2.0 SPS05

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version		HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version			
	RHEL 7.3	Supported	Not Supported
	RHEL 7.4	Supported	Not Supported
	RHEL 7.5	Supported	Not Supported
	RHEL 7.6	Supported	Supported
	RHEL 8.1	Not Supported	Supported

SUSE Linux

Supported DBMS Version		HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
	Supported OS Version		
	SUSE 12 SP3	Supported	Not Supported
	SUSE 15 SP2	Not Supported	Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1	HANA 4.1.2	HANA 4.2.0	HANA 4.3.0	HANA 4.3.1	HANA 4.3.2
6.0.0.x	Yes										
6.0.1.x	Yes										
6.0.2.x	Yes										
6.0.3.x	Yes										
6.0.4.x	Yes										

6.0.5.x	Yes										
6.0.6.x	Yes										
6.0.7.x	Yes										
6.0.8.x	Yes										
6.0.9.x	Yes										

Unsupported Versions and Features

- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Adding or removing the Hana services on the source tenant after dSource creation is not supported in case of Delphix-initiated backups. In case of customer provided or CommVault, full backups will have to be created after adding or removing any services.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- Using same HANA instance with different plugin version on different Delphix Engines is not supported.
- Changing the Backup Path after initial dSource creation is not supported.
- V2P is not supported with this release.
- Third-party backup tools other than CommVault are not supported.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.3.2
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0

Your Version	Recommended upgrade path to 4.3.2
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
4.0.0	Upgrade to version 4.1.1 and then to 4.1.2 and then to 4.2.0
4.1.1	Upgrade directly to version 4.1.2 and then to 4.2.0
4.1.2	Upgrade directly to version 4.2.0
4.2.0	Upgrade directly to versions 4.3.1 and 4.3.2

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None
DLPX-44359	Job cancellation may fail.	None
CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1/4.1.2/4.2.0 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.

Fixed Issues

The following issues are fixed in this release:

Delphix Ticket Number	Summary
CTCHANA-1666	Backporting the password change through hooks issue to the master branch
CTCHANA-1660	remove dependency on /etc/SuSE-release for suse o/s platforms

HANA 2.0 Plugin Version 4.3.1 Release Notes

This is the release of the HANA plugin version 4.3.1.

New Features

No new features.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS03
- HANA Standard and Enterprise Edition 2.0 SPS04
- HANA Standard and Enterprise Edition 2.0 SPS05

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version		HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version				
	RHEL 7.3	Supported	Supported	Not Supported
	RHEL 7.4	Supported	Supported	Not Supported
	RHEL 7.5	Supported	Supported	Not Supported
	RHEL 7.6	Supported	Supported	Not Supported
	RHEL 8.1	Not Supported	Not Supported	Supported

SUSE Linux

Supported DBMS Version		HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version				
	SUSE 12 SP3	Supported	Supported	Not Supported
	SUSE 12 SP4	Supported	Supported	Not Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1	HANA 4.1.2	HANA 4.2.0	HANA 4.3.0	HANA 4.3.1
6.0.0.x	Yes									
6.0.1.x	Yes									
6.0.2.x	Yes									
6.0.3.x	Yes									
6.0.4.x	Yes									
6.0.5.x	Yes									
6.0.6.x	Yes									
6.0.7.x	Yes									
6.0.8.x	Yes									

Unsupported Versions and Features

- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.

- Adding or removing the Hana services on the source tenant after dSource creation is not supported in case of Delphix-initiated backups. In case of customer provided or CommVault, full backups will have to be created after adding or removing any services.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- Using same HANA instance with different plugin version on different Delphix Engines is not supported.
- Changing the Backup Path after initial dSource creation is not supported.
- V2P is not supported with this release.
- Third-party backup tools other than CommVault are not supported.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.3.1
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
4.0.0	Upgrade to version 4.1.1 and then to 4.1.2 and then to 4.2.0
4.1.1	Upgrade directly to version 4.1.2 and then to 4.2.0
4.1.2	Upgrade directly to version 4.2.0
4.2.0	Upgrade directly to version 4.3.1

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None.
DLPX-44359	Job cancellation may fail.	None.

CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1/4.1.2/4.2.0 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.
-------------	--	--

HANA 2.0 Plugin Version 4.3.0 Release Notes

This is the release of the HANA plugin version 4.3.0.

New Features

Below features are being provided:-

1. **Non-ingestion of backups into Delphix:** This means that database backups will not be transferred into the Delphix Engine during any of the operation. This will be driven on the basis of a checkbox that can be selected by the customer during dSource creation. More details have been provided [here](#).
2. Support provided for log backups along with the data backups (only when non-ingestion has been chosen during dSource creation). More details have been provided [here](#).
3. Ability to use Delphix provided storage for all virtualized tenant services. This is valid even for the user-created services outside the plugin.
4. Multiple HANA installations on the same VM will also be supported.
5. Support for Self-Service (Jetstream) has been certified with this release.
6. Certification of plugin support with HANA Standard and Enterprise Edition 2.0 SP05 on RHEL 8.1.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS03
- HANA Standard and Enterprise Edition 2.0 SPS04
- HANA Standard and Enterprise Edition 2.0 SPS05

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
Supported OS Version			

	RHEL 7.3	Supported	Supported	Not Supported
	RHEL 7.4	Supported	Supported	Not Supported
	RHEL 7.5	Supported	Supported	Not Supported
	RHEL 7.6	Supported	Supported	Not Supported
	RHEL 8.1	Not Supported	Not Supported	Supported

SUSE Linux

Supported DBMS Version		HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
	Supported OS Version			
	SUSE 12 SP3	Supported	Supported	Not Supported
	SUSE 12 SP4	Supported	Supported	Not Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1	HANA 4.1.2	HANA 4.2.0	HANA 4.3.0
6.0.0.x	Yes								
6.0.1.x	Yes								
6.0.2.x	Yes								

6.0.3.x	Yes								
6.0.4.x	Yes								
6.0.5.0	Yes								

Unsupported Versions and Features

- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Adding or removing the Hana services on the source tenant after dSource creation is not supported in case of Delphix-initiated backups. In case of customer provided or CommVault, full backups will have to be created after adding or removing any services.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- Using same HANA instance with different plugin version on different Delphix Engines is not supported.
- Changing the Backup Path after initial dSource creation is not supported.
- V2P is not supported with this release.
- Third-party backup tools other than CommVault are not supported.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.1.2
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
4.0.0	Upgrade to version 4.1.1 and then to 4.1.2 and then to 4.2.0
4.1.1	Upgrade directly to version 4.1.2 and then to 4.2.0
4.1.2	Upgrade directly to version 4.2.0
4.2.0	Upgrade directly to version 4.3.0

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None.
DLPX-44359	Job cancellation may fail.	None.
CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1/4.1.2/4.2.0 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.

HANA 2.0 Plugin Version 4.2.0 Release Notes

This is the release of the HANA plugin version 4.2.0.

New Features

With this release, we are including the support for `basepath_databackup` parameter. Backups created via plugin will now be created on the PATH specified by this parameter. This parameter can be defined at the system and at the tenant level. If this parameter has not been specified then default location will be used for creating the backups.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualisation Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS03
- HANA Standard and Enterprise Edition 2.0 SPS04

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version		HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)
	Supported OS Version		
	RHEL 7.3	Supported	Supported
	RHEL 7.4	Supported	Supported
	RHEL 7.5	Supported	Supported
	RHEL 7.6	Supported	Supported

SUSE Linux

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)
	Supported OS Version	
SUSE 12 SP3	Supported	Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1	HANA 4.1.2	HANA 4.2.0
6.0.0.x	Yes							
6.0.1.x	Yes							
6.0.2.x	Yes							
6.0.3.x	Yes							
6.0.4.x	Yes							

Unsupported Versions and Features

- The HANA 2.0 Plugin does not support multiple HANA installations on the same host.
- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Parallel or concurrent VDB creation against a single target host is not supported.
- Third-party backup tools other than Commvault are not supported.
- Scale-out source and target environments are not supported.
- V2P is not supported with this release.
- Only child VDBs should be made available to Self-Service users.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.1.2
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1 and then to 4.1.2 and then to 4.2.0
4.0.0	Upgrade to version 4.1.1 and then to 4.1.2 and then to 4.2.0
4.1.1	Upgrade directly to version 4.1.2 and then to 4.2.0
4.1.2	Upgrade directly to version 4.2.0

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None.
DLPX-44359	Job cancellation may fail.	None.
CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1/4.1.2/4.2.0 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.

HANA 2.0 Plugin Version 4.1.2 Release Notes

This is the release of the HANA plugin version 4.1.2.

New Features

No new features.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualisation Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS03
- HANA Standard and Enterprise Edition 2.0 SPS04

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)
Supported OS Version		
RHEL 7.3	Supported	Supported
RHEL 7.4	Supported	Supported
RHEL 7.5	Supported	Supported
RHEL 7.6	Supported	Supported

SUSE Linux

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)

Supported OS Version		
SUSE 12 SP3	Supported	Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1	HANA 4.1.2
6.0.0.x	Yes						
6.0.1.x	Yes						
6.0.2.x	Yes						
6.0.3.x	Yes						

Unsupported Versions and Features

- The HANA 2.0 Plugin does not support multiple HANA installations on the same host.
- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Parallel or concurrent VDB creation against a single target host is not supported.
- Third-party backup tools other than Commvault are not supported.
- Scale-out source and target environments are not supported.
- V2P is not supported with this release.
- Only child VDBs should be made available to Self-Service users.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.1.2
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1 and then to 4.1.2
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1 and then to 4.1.2
4.0.0	Upgrade to version 4.1.1 and then to 4.1.2

Your Version	Recommended upgrade path to 4.1.2
4.1.1	Upgrade directly to version 4.1.2

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None.
DLPX-44359	Job cancellation may fail.	None.
CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1/4.1.2 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.

HANA 2.0 Plugin Version 4.1.1 Release Notes

This is the release of the HANA plugin version 4.1.1.

New Features

No new features.

Plugin Installation

Plugin to support HANA virtualization are not automatically included in the Virtualisation Engine deployment and are installed by Delphix Professional Services or Technical Support. A support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported DBMS Versions

- HANA Standard and Enterprise Edition 2.0 SPS03
- HANA Standard and Enterprise Edition 2.0 SPS04

Supported Operating Systems

- Red Hat Enterprise Linux (RHEL)
- S.u.S.E Linux

Red Hat Enterprise Linux (RHEL)

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)
Supported OS Version		
RHEL 7.3	Supported	Supported
RHEL 7.4	Supported	Supported
RHEL 7.5	Supported	Supported
RHEL 7.6	Supported	Supported

SUSE Linux

Supported DBMS Version	HANA 2.0 SPS 03 (Standard and Enterprise Editions)	HANA 2.0 SPS 04 (Standard and Enterprise Editions)

Supported OS Version		
SUSE 12 SP3	Supported	Supported

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	HANA 3.4.5	HANA 3.5.0	HANA 3.5.1	HANA 4.0.0	HANA 4.1.0	HANA 4.1.1
6.0.0.x	Yes	Yes	Yes	Yes	Yes	Yes
6.0.1.x	Yes	Yes	Yes	Yes	Yes	Yes
6.0.2.x	Yes	Yes	Yes	Yes	Yes	Yes

Unsupported Versions and Features

- The HANA 2.0 Plugin does not support multiple HANA installations on the same host.
- Multiple duplicate services are not supported in this release.
- Point in Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Parallel or concurrent VDB creation against a single target host is not supported.
- Third-party backup tools other than Commvault are not supported.
- Scale-out source and target environments are not supported.
- V2P is not supported with this release.
- Only child VDBs should be made available to Self-Service users.
- Script server HANA service is not supported with this release.

Plugin Upgrade Path

Your Version	Recommended upgrade path to 4.1.1
3.4.5	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1
3.5.0	Upgrade to version 3.5.1 and then upgrade to 4.0.0 and then to 4.1.1
3.5.1	Upgrade to version 4.0.0 and then to 4.1.1
4.0.0	Upgrade directly to version 4.1.1

Known Issues

The following issues are known to exist in this release:

Delphix Ticket Number	Summary	Workaround/Comments
HANA-821	HANA SQL Port value should remain consistent for the life of VDB.	None.
DLPX-44359	Job cancellation may fail.	None.
CTCHANA-930	Upgrade from 3.4.5 to 3.5.0/3.5.1/4.0.0/4.1.1 needs a manual deletion of the existing tenant databases for disable/enable to work.	Drop the tenant database manually before enabling the VDB.

PostgreSQL Release Notes

- [PostgreSQL 1.5.1 Release Notes](#)
- [PostgreSQL 1.5.0 Release Notes](#)
- [PostgreSQL 1.4.4 Release Notes](#)
- [PostgreSQL 1.4.3 Release Notes](#)
- [PostgreSQL 1.4.2 Release Notes](#)
- [PostgreSQL 1.4.1 Release Notes](#)
- [PostgreSQL 1.4.0 Release Notes](#)

PostgreSQL 1.5.1 Release Notes

This release includes the Streaming Replication Monitoring feature that monitors the streaming replication by reading the Postgres database logs every minute.

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	■
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	■
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where **x** represents the minor version.

 • All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.

• Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

 Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.5.1 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x, 12.x, and 13.x customers

Path B - Existing PostgreSQL customers with different Plugin versions:

Your Version	Recommended Upgrade Path to 1.5.1
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1, 1.4.2, 1.4.3, or 1.4.4 and follow the path below.

1.4.1	Upgrade directly to 1.5.0 or the latest version 1.5.1.
1.4.2	
1.4.3	

For PostgreSQL customers with plugin version 1.0.1 or 1.0.2 interested in the External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0/1.5.1

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0/1.5.1

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0/1.5.1

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0/1.5.1

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using the "custom Configuration" option on the UI :

- Source IP Address
- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolved on the status tab of dSource.

Fixed Issues

PostgreSQL version 1.5.1 contains the following bug fixes and enhancements:

Ticket Number	Summary
POST-418	Replication monitoring - Lua plugin

PostgreSQL 1.5.0 Release Notes

This release includes some bug fixes and enhancements (listed in this document towards the end).

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	■
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	■
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where x represents the minor version.

 • All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.

• Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

 Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.5.0 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x, 12.x, and 13.x customers

Path B - Existing PostgreSQL customers with different Plugin versions:

Your Version	Recommended Upgrade Path to 1.5.0
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1, 1.4.2, 1.4.3, or 1.4.4 and follow the path below.

1.4.1	Upgrade directly to the latest version, 1.5.0.
1.4.2	
1.4.3	

For PostgreSQL customers with plugin version 1.0.1 or 1.0.2 interested in the External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4 → 1.5.0

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using the "custom Configuration" option on the UI :

- Source IP Address
- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolved on the status tab of dSource.

Fixed Issues

PostgreSQL version 1.5.0 contains the following bug fixes and enhancements:

Ticket Number	Summary
POST-202	If the port is not present in postgresql.conf on the source, it's not provided in staging and VDBs
POST-283	Security enhancement - Plugin overrides listen_addresses parameter in postgresql.conf to "localhost" on staging & target server
POST-292	dSource configuration parameters don't accept empty values
POST-295	The staging server is created using a copy of the production parameters file
POST-297	Postgres Plugin Should check if the toolkit directory is accessible by the OS user

POST-298	primary_conninfo' parameter is not getting commented out in vdb flow
POST-302	Allow users to set empty values for a parameter in a configuration file
POST-312	delphix_postgres_debug.log is too large to manage
POST-316	Propagation of Hotfix 1112 code changes to plugin 1.5.0 release

PostgreSQL 1.4.4 Release Notes

Welcome to the latest release of the Delphix Engine database virtualization system.

New Features

No new features are included in PostgreSQL plugin version 1.4.4. This release includes a bug fix (listed in this document towards the end).

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where **x** represents the minor version.

 • All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.

• Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

 Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.4.4 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x, 12.x, and 13.x customers

Path B - Existing PostgreSQL customers with different Plugin versions:

Your Version	Recommended Upgrade Path to 1.4.4
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1 or 1.4.2, 1.4.3, or 1.4.4 the latest plugin for PostgreSQL.

For PostgreSQL customers with plugin version 1.0.1 or 1.0.2 interested in the External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3/1.4.4

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using the "custom Configuration" option on the UI :

- Source IP Address
- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolved on the status tab of dSource.

Fixed Issues

PostgreSQL version 1.4.4 contains the following bug fix:

Key	Summary
POST-291	Error Handling for Authentication Failure in <code>pg_basebackup</code> command

PostgreSQL 1.4.3 Release Notes

Welcome to the latest release of the Delphix Engine database virtualization system.

New Features

No new features are included in PostgreSQL plugin version 1.4.3. This release includes a bug fix (listed in this document towards the end).

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	■
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	■
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where **x** represents the minor version.

-  All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.
- Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

-  The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

-  Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.4.3 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x, 12.x, and 13.x customers

Path B - Existing PostgreSQL customers with different Plugin versions:

Your Version	Recommended Upgrade Path to 1.4.3
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1 or 1.4.2 or 1.4.3 the latest plugin for PostgreSQL.

For PostgreSQL customers with plugin version 1.0.1 or 1.0.2 interested in the External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1/1.4.2/1.4.3

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using the "custom Configuration" option on the UI :

- Source IP Address
- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolve on the status tab of dSource.

Fixed Issues

PostgreSQL version 1.4.3 contains the following bug fix:

Key	Summary
POST-284	The plugin doesn't check if the staging Postgres database is in a stable state or not

PostgreSQL 1.4.2 Release Notes

Welcome to the latest release of the Delphix Engine database virtualization system.

New Features

No new features included in PostgreSQL plugin version 1.4.2. This release consists of only bug fixes (listed in this document towards the end).

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported Versions

PostgreSQL Standard Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

PostgreSQL EDB Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

Supported Operating Systems

RHEL/CentOS/SLES - PostgreSQL Open Source & EDB

Supported OS Version	Supported DBMS Version		
	PostgreSQL 9.4.x (Open Source & EDB)	PostgreSQL 9.5.x (Open Source & EDB)	PostgreSQL 9.6.x (Open Source & EDB)
RHEL/CentOS 7.3	Supported	Supported	Supported
RHEL/CentOS 7.4	Supported	Supported	Supported

RHEL/CentOS 7.5	Supported	Supported	Supported
RHEL/CentOS 7.6	Supported	Supported	Supported
RHEL/CentOS 7.8	Not Supported	Not Supported	Not Supported
SLES 11	Supported	Not Supported	Not Supported
SLES 12	Supported	Supported	Supported

where **x** is the minor version of the PostgreSQL Database.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

Delphix Engine	PostgreSQL 1.0.1	PostgreSQL 1.0.2	PostgreSQL 1.1.0	PostgreSQL 1.2.0	PostgreSQL 1.3.0	PostgreSQL 1.3.1	PostgreSQL 1.3.2	PostgreSQL 1.4.0	PostgreSQL 1.4.1	PostgreSQL 1.4.2
6.0.0.x	No	No	Yes							
6.0.1.x	No	No	Yes							
6.0.2.x	No	No	Yes							
6.0.3.x	No	No	Yes							
6.0.4.x	No	No	Yes							
6.0.5.x	No	No	Yes							
6.0.6.x	No	No	Yes							
6.0.7.x	No	No	Yes							
6.0.8.x	No	No	Yes							
6.0.9.x	No	No	Yes							

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.4.2 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x, and 12.x customers

Path B - Existing PostgreSQL customers with different Plugin version:

Your Version	Recommended Upgrade Path to 1.4.2
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1 or 1.4.2 the latest plugin for PostgreSQL.

For PostgreSQL customers with plugin version 1.0.1 or 1.0.2 interested in the External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1/1.4.2

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1/1.4.2

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using the "custom Configuration" option on the UI :

- Source IP Address
- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolve on the status tab of dSource.

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e 9.2, 9.3 (including the minor versions) are not supported.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - **Unlogged Tables:** All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - **Point In Time recovery:** Currently, the plugin doesn't support Point in Time recovery of the PostgreSQL database.

Fixed Issues

PostgreSQL version 1.4.2 contains the following bug fixes:

Key	Summary
POST-274	`data_directory` parameter in the postgresql.conf file is causing failure in dSource and VDB flow
POST-275	Linking Fails if the parameter which needs to be configured is commented through multiple `#`
POST-276	Linking Fails if the mandatory parameter required to setup replication is commented in config file

PostgreSQL 1.4.1 Release Notes

Welcome to the latest release of the Delphix Engine database virtualization system.

New Features

No new features included in PostgreSQL Plugin version 1.4.1. This release consists of only bug fixes (listed in this document towards the end).

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported Versions

PostgreSQL Standard Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

PostgreSQL EDB Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

Supported Operating Systems

RHEL/CentOS/SLES - PostgreSQL Open Source & EDB

Supported DBMS Version	PostgreSQL 9.4.x (Open Source & EDB)	PostgreSQL 9.5.x (Open Source & EDB)	PostgreSQL 9.6.x (Open Source & EDB)	PostgreSQL 10.x (Open Source & EDB)	PostgreSQL 11.x (Open Source & EDB)	PostgreSQL 12.x (Open Source & EDB)
Supported OS Version						

RHEL/ CentOS 7.3	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.4	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.5	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.6	Supported	Supported	Supported	Supported	Supported	Supported
SLES 11	Supported	Not Supported				
SLES 12	Supported	Supported	Supported	Supported	Supported	Not Supported

where **x** is the minor version of the PostgreSQL Database.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

Delphi x Engine	Postgre SQL 1.0.1	Postgre SQL 1.0.2	Postgre SQL 1.1.0	Postgre SQL 1.2.0	Postgre SQL 1.3.0	Postgre SQL 1.3.1	Postgre SQL 1.3.2	Postgre SQL 1.4.0	Postgre SQL 1.4.1
6.0.0.x	No	No	Yes						
6.0.1.x	No	No	Yes						
6.0.2.x	No	No	Yes						
6.0.3.x	No	No	Yes						
6.0.4.x	No	No	Yes						
6.0.5.x	No	No	Yes						
6.0.6.x	No	No	Yes						

6.0.7.x	No	No	Yes						
6.0.8.x	No	No	Yes						
6.0.9.x	No	No	Yes						

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.4.1 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x, 10.x, 11.x, and 12.x customers

Path B - Existing PostgreSQL customers with different Plugin version:

Your Version	Recommended Upgrade Path to 1.4.1
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.1, the latest plugin for PostgreSQL.

For PostgreSQL customers with Plugin version 1.0.1 or 1.0.2 interested in External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.1

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.1

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using "custom Configuration" option on the UI :

- Source IP Address

- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolve on the status tab of dSource.

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e 9.2, 9.3 (including the minor versions) are not supported.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - **Unlogged Tables:** All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - **Point In Time recovery:** Currently, the Plugin doesn't support Point in Time recovery of PostgreSQL database

Fixed Issues

PostgreSQL version 1.4.1 contains the following bug fixes:

Key	Summary
POST-269	Postgres DB_SYNC fails - ESCL-3179
POST-270	Plugin should handle postgresql.auto.conf entries during dsource creation

PostgreSQL 1.4.0 Release Notes

Welcome to the latest release of the Delphix Engine database virtualization system.

New Features

PostgreSQL Plugin version 1.4.0 contains the following certifications :

Operating System - RHEL/CentOS 7.8 version Support with PostgreSQL 12.x.

Plugin Installation

Plugins to support PostgreSQL virtualization are not automatically included in the Virtualization Engine deployment and are installed by Delphix Professional Services or Technical Support. A Support case can be opened requesting installation by following this [link](#).

Migration and Compatibility

Supported Versions

PostgreSQL Standard Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

PostgreSQL EDB Versions:

- PostgreSQL 9.4.x
- PostgreSQL 9.5.x
- PostgreSQL 9.6.x
- PostgreSQL 10.x
- PostgreSQL 11.x
- PostgreSQL 12.x

Supported Operating Systems

RHEL/CentOS/SLES - PostgreSQL Open Source & EDB

Supported DBMS Version	PostgreSQL 9.4.x (Open Source & EDB)	PostgreSQL 9.5.x (Open Source & EDB)	PostgreSQL 9.6.x (Open Source & EDB)	PostgreSQL 10.x (Open Source & EDB)	PostgreSQL 11.x (Open Source & EDB)	PostgreSQL 12.x (Open Source & EDB)
Supported OS Version						

RHEL/ CentOS 7.3	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.4	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.5	Supported	Supported	Supported	Supported	Supported	Supported
RHEL/ CentOS 7.6	Supported	Supported	Supported	Supported	Supported	Supported
SLES 11	Supported	Not Supported	Not Supported	Not Supported	Not Supported	Not Supported
SLES 12	Supported	Supported	Supported	Supported	Supported	Not Supported

where **x** is the minor version of the PostgreSQL Database.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

Delphix Engine	PostgreS QL 1.0.1	PostgreS QL 1.0.2	PostgreS QL 1.1.0	PostgreS QL 1.2.0	PostgreS QL 1.3.0	PostgreS QL 1.3.1	PostgreS QL 1.3.2	PostgreS QL 1.4.0
6.0.0.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.1.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.2.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.3.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.4.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.5.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.6.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes

6.0.7.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.8.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes
6.0.9.x	No	No	Yes	Yes	Yes	Yes	Yes	Yes

Plugin Upgrade Path

Path A - New Plugin installation for PostgreSQL customers:

- Upload 1.4.0 Plugin for PostgreSQL 9.4.x, 9.5.x, 9.6.x,10.x, 11.x and 12.x customers

Path B - Existing PostgreSQL customers with different Plugin version:

<i>Your Version</i>	<i>Recommended Upgrade Path to 1.4.0</i>
1.0.1 1.0.2	Upgrade to 1.1.0 and follow the upgrade path below.
1.1.0	Upgrade to 1.2.0 and follow the upgrade path below.
1.2.0	Upgrade to 1.3.0 and follow the upgrade path below.
1.3.0	Upgrade to 1.4.0, the latest plugin for PostgreSQL.

For PostgreSQL customers with Plugin version 1.0.1 or 1.0.2 interested in External Backups feature:

Plugin upgrade path: 1.0.1 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0

Plugin upgrade path 1.0.2 → 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0

For PostgreSQL customers with Plugin version 1.1.0

Plugin upgrade path: 1.1.0 → 1.2.0 → 1.3.0 → 1.4.0

For PostgreSQL customer with Plugin version 1.2.0

Plugin upgrade path: 1.2.0 → 1.3.0 → 1.4.0

Upgrading Existing Objects

In order to upgrade the existing objects with 1.1.0 (i.e 1.0.1 → 1.1.0 or 1.0.2 → 1.1.0), just after the upgrade, it is mandatory to update below two parameters of dSource using "custom Configuration" option on the UI :

- Source IP Address

- Source Port

A critical fault will show up on dSource just after the upgrade. It is because the engine is not able to find the source config from which dSource was created. As we have moved to manual discovery in this release, we need to mark this fault as resolve on the status tab of dSource.

Unsupported PostgreSQL Versions and Features

PostgreSQL versions below 9.4 i.e 9.2, 9.3 (including the minor versions) are not supported.

Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.

Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of PostgreSQL database

Overview of Data Virtualization



Unknown macro: 'navbox'

Delphix Engine Overview

What is Delphix?

Delphix is a data management platform that provides the ability to securely copy and share datasets. Using virtualization, you will ingest your data sources and create virtual data copies, which are full read-write capable database instances that use a small fraction of the resources a normal database copy would require.

Overview

The Delphix Engine links to source physical databases via standard APIs and asks the source databases to send copies of their entire file and log blocks to it. The copy of the source database stored in the Delphix Engine, along with all incremental updates, is referred to as the dSource in Delphix terminology.

After the initial loading, the Delphix Engine maintains synchronization with source databases based on a user-defined policy. Once linked, Delphix maintains a Timeflow of the source database - a record of snapshots and log changes. From any time within that Timeflow, a virtual database (referred to in Delphix terminology as a VDB) can be instantly provisioned from the Delphix Engine. VDBs are served from the shared storage footprint of the dSource database Timeflow, so no additional storage is required.

Multiple VDBs can be provisioned from any point in time in a Timeflow, down to the second. Once provisioned, a VDB is an independent, read-write database, and changes made to the VDB by users or applications are written to new, compressed blocks in Delphix storage. VDBs can be provisioned from other VDBs, and the data within VDBs refreshed from its parent VDB or dSource.

Related Topics

- [Getting Started with Data Sources](#)
- [Glossary of Major Delphix Concepts](#)

Getting Started with Data Sources

Delphix enables agile data management, but our users first need to learn how to use our platform. Learn more about what data management looks like and how Delphix fits into the bigger picture. The following sections link to the Getting Started section of our Datasets category. This category also has explanations of other components of the Delphix platform, such as environments and policies.

Getting Started Outline

- [Managing Environments and Hosts](#)
- [Managing Data Sources and Syncing Data](#)
- [Provisioning and Managing Virtual Databases](#)
- [Hook Scripts for Automation and Customization](#)
- [Policies for Scheduled Jobs](#)

Managing Data Sources and Syncing Data

This section explains how to manage your data sources and create dSources. It hopes to answer questions such as: when you link a database, what happens? How does Delphix stay in sync with changes to a data source, and is there any impact to that source?

Learn about how we link to your data sources to enable fast and secure data management.

Provisioning and Managing Virtual Databases

This section details the virtual database (VDB) functionality and explains how you can create, manage, and utilize these data objects. This section will address questions such as: what is a virtual database, what makes them virtual? What does it mean to provision?

Learn all about the objects we create and refer to as virtual databases or VDBs.

Related Topics

- [Managing Environments and Hosts](#)
- [Managing Data Sources and Syncing Data](#)
- [Provisioning and Managing Virtual Databases](#)
- [Hook Scripts for Automation and Customization](#)
- [Policies for Scheduled Jobs](#)

Delphix Product Information

Product Overview

The Delphix Engine is delivered and maintained as a closed virtual software appliance. Like all virtual appliances, the Delphix Engine is a tightly integrated combination of a special-purpose operating system and business logic. A single engine can be configured for data virtualization or data masking.

The product is delivered as a closed appliance because there are dependencies between software components within the virtual appliance, which require end-to-end testing. As such, we do not provide administrative access to the operating system for any reason, including installing software, make customizations, or perform security scans. More details about the administrative model will be provided in later sections of this document.

Product Packaging

The DDDP product is delivered and maintained as a virtual software appliance. The mode of product delivery is dependent on the hosting hypervisor platform described in the table below:

Hypervisor Platform	Delivery Vehicle	Notes
VMware ESXi	OVA image	
Amazon AWS EC2	AMI image	Manual deployment to EC2
Amazon AWS EC2	Guest Machine Appliance	Amazon Marketplace
Microsoft Azure	Guest Machine Appliance	Azure Marketplace
Google Cloud Platform	Guest Machine Appliance	
Hyper V	Hyper image	

Product Updates, Upgrades, and Versions

Regardless of the initial packaging used to deploy the Delphix Engine, updates are supplied as a single upgrade image of the new release, and the same image can be used for any prior release from which an upgrade is supported. This upgrade image delivers a completely new appliance, including both the operating system and business logic components.

The upgrade process retains prior configuration and customer data such that no data or configuration is lost during the upgrade process. The upgrade process retains a copy of the previous version of the components for automatic recovery in case an upgrade fails. Installed versions older than the prior version are automatically deleted during the upgrade.

Delphix characterizes software as Major Release, Minor Release, and Maintenance Release in the Delphix Support Policies. These release types indicate the level of feature addition and change. For example, a Major Release may introduce significant new features, interface changes, and many bug fixes. A Maintenance Release or Patch Release

may deliver only a small number of bug fixes and no feature additions. In any case, each release delivers a complete upgrade image of the appliance.

There is no component patching of the Delphix Engine; fixes are delivered in new versions of the software as a new software appliance. There is no management of patches involved, and each Delphix Engine version is a consistent, tested virtual appliance.

Product Administration

Administration of the Delphix Engine is effected through product interfaces only. These interfaces provide for the proper configuration and testing of customer infrastructure components, such as network addresses, storage, Domain Name Service (DNS) servers, authentication servers (LDAP), etc. The interfaces also control the business logic and control of the overall platform, including how customer data is used and provisioned by the system.

Although the special-purpose operating system may be accessed by Delphix Support and Engineering personnel for the purpose of diagnostics and problem remediation, there are no customer-accessible interfaces at the operating system level. Customers are not provided access to the underlying operating system nor can any custom software be installed on the appliance.

Product Customization

The product has several endpoints allowing customization for improved integration with customer environments, local business workflow requirements, and alternative data sources.

Interface	Functional Area	Description
Data Plugins	Data Virtualization	Delivered by Delphix Services or Integration partners. These plugins allow for supporting additional data types including both structured and unstructured data.
Privilege Elevation Profiles	Data Virtualization	Delivered by Delphix Services or Integration Partners. These customizations allow for the use of privilege mechanisms other than sudo on Linux and Unix target environments. Sudo is the product default.
Hook Scripts	Data Virtualization	These customer-managed scripts allow for custom business logic to be applied to Oracle and SQL Server data sources and virtual databases. The scripts are not integrated into the appliance but are referenced and invoked by the product during data operations.
Custom Algorithms	Data Masking	Delivered by Delphix Services or Integration Partners. Custom algorithms provide specialized data transformations to secure or anonymize sensitive data.

In addition to these endpoints, Delphix provides a robust set of application programming interfaces (APIs) that enable business automation and fully integrated data operations into client workflows.

Related Links

- [Delphix Product Security](#)

Glossary of Major Delphix Concepts

This glossary is your guide to exploring Delphix terms and definitions.

Products

Term	Explanation
Delphix Virtualization	The Delphix product to deliver data on-demand to application developers and testers.
Delphix Masking	The Delphix product to identify sensitive data and replace it with realistic but fictitious data.

Functional Product Components

Term	Explanation
Delphix Management	The user interface for Delphix administrators to configure the product
Delphix Self-Service	The user interface designed specifically for project teams, application developers and testers.

Delphix Virtualization Engine Terms

Delphix Virtualization Concepts

Term	Explanation
Blocks or data blocks	Relational Database Management System (RDBMS) databases are based on files containing data blocks on disk. These are then backed up and restored with traditional database tools or virtualized with Delphix.
Dataset	A generic term for data in Delphix Virtualization. This can refer to sources or copies of data across any file/database.

Data source	<p>The set of data, typically a database, that you would like to create copies of. Data sources serve as the base from which you will create virtual copies. The source may be a database, a cluster of databases, or a file system.</p> <p>Not to be confused with a dSource, which is a virtualized, compressed duplicate of this database. (See below.)</p>
dSource	<p>A database that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools. Because dSources are simply sourced data, you must provision a VDB in order to distribute/clone/test the data being pulled in. VDBs can also later be refreshed from the same or other points in time synced from the dSource.</p>
Delphix Connector	<p>A service that runs on a Windows host and enables the communication between the Delphix Virtualization Engine and the Windows Target Environment where it is installed.</p>
Environment	<p>The server and software required to run a data set. For example, a Linux system running Postgres. This may either be an individual instance, or a cluster like Oracle RAC.</p> <p>Environments can either be a source (where data comes from), staging (where data are prepared/masked) or target (where data are delivered and used by developers and testers).</p>
Hooks	<p>Mechanisms by which Delphix can run scripts or call external processes during common Delphix operations. For example, running a particular script during every VDB provision.</p>
HostChecker	<p>A standalone program that validates that environments are configured correctly before the Delphix Virtualization Engine uses them for syncing from data sources or provisioning VDBs.</p> <p>HostChecker should run before adding any Environment to your Delphix Virtualization Engine. It is available to download from the HostChecker subfolder at download.delphix.com.</p>
Replication	<p>Replication enables creating a copy of specific data sets on a second Delphix Virtualization Engine. This is used both for disaster recovery purposes as well as scale-out. Replication is configured on a source engine and the data are moved to a secondary engine. The source engine then sends incremental updates manually or according to a schedule.</p>

Snapshots	<p>Snapshots represent the state of a dataset at a specific moment in time. Snapshots are created from policies or are generated by manual creation. Snapshots allow you to choose a point in time from which to provision, refresh, or rollback.</p> <p>If you have LogSync enabled, you can provision refresh or rollback from a point in time between the snapshots endpoints.</p> <p>Note: Point in time provisioning is also dependent on log retention.</p>
SnapSync	<p>The standard process for importing data from a source into the Delphix Virtualization Engine. An initial SnapSync is performed to create a dSource on the Delphix Virtualization Engine. Incremental SnapSyncs are performed to provide additional points in time to the dSource on the Delphix Virtualization Engine.</p>
LogSync	<p>This feature enables the ingestion and retention of more granular (log-based) source change data. This more granular change data allows for VDB point-in-time provision, refresh, or rollback.</p>
Validated Sync	<p>The process that runs on a staging database within a Staging Environment executes database cleanup and preparation. This is recommended and results in faster and more predictable provisions and refreshes, and which executes either before a snapshot is taken (SQL Server) or after a snapshot is taken (Oracle).</p>
Timeflow	<p>The collection of snapshots for a particular data set. Virtual databases can be provisioned from any snapshots in a Timeflow.</p>
Unstructured files	<p>Data stored in a filesystem that is not part of a database Unstructured files can consist of anything from a simple directory to the root of a complex application like Oracle E-Business Suite. Like with other data types, you can configure a dSource to sync periodically with a set of unstructured files external to the Delphix Virtualization Engine.</p>
Virtual Database (VDB)	<p>A database provisioned from either a dSource or another VDB which is a full read/write copy of the source data. A VDB is created and managed by the Delphix Virtualization Engine.</p>
Virtual dataset	<p>A comprehensive term that includes VDBs and virtualized files.</p>
V2P	<p>Shorthand for “Virtual-to-Physical Provisioning”. This refers to the process of moving a dataset from Delphix back to a full-size database. This can be used in DR purposes or as part of a final performance test, matching the production systems.</p>

TDE Definitions

The following terms are used by Delphix and are summarized here for clarity.

Term	Definition
Artifact directory	Directory on the target system (not on Delphix storage) which stores keys needed to support Delphix workflows on TDE-enabled vPDBs. It is located under the toolkit directory.
Auxiliary container database (CDB)	Provisioning an Oracle vPDB requires running recovery to bring the snapshotted datafiles into a consistent state. This needs to be done in the context of a container database, which is created on the target system. After the recovery is complete, the vPDB is unplugged and plugged into the target container, and the auxiliary container is deleted.
Exported keyfile	File located on the target Oracle host which contains keys that have been exported from the Keystore. It is encrypted with a secret that is specified when it is exported. The exported keyfile itself cannot be used as a Keystore, but its contents can be imported into a new Keystore.
Key rotation	Process for changing the master encryption key in the keystore via ADMINISTER KEY MANAGEMENT SET KEY. This does not remove the original key, rather it adds a new key to the wallet and future data will be encrypted with the new key.
Keystore/wallet	File found on the Oracle host which stores the keys used to encrypt and decrypt the internal table keys in a database. Every keystore has a password that is set when it is first created and must be supplied for operations on it.
Parent Keystore	Keystore with the keys used to encrypt the dSource PDB files.
Target Keystore	Keystore for the target CDB into which the TDE-encrypted vPDB is plugged.

Dataset Operations

The terms below describe actions you can perform on a Delphix Virtualization Engine.

Term	Explanation
------	-------------

Link	The process of establishing a relationship between a data source and the Delphix Virtualization Engine. After linking a data source, the Delphix Virtualization Engine can import data periodically and manage it as it evolves over time. In the user interface, this accomplished via "Add dSource."
Mask	Masking replaces sensitive data with fictitious data in non-prod environments (such as VDBs). It provides realistic data with which to work while reducing security risks. For more details about masking, see Masking Terms below.
Migrating a VDB	Moving a VDB to a new Target Environment.
Provision	Create a new VDB from a dSource or VDB.
Refresh	Refreshing a VDB will re-provision it from the dSource. As with the normal provisioning process, you can choose to refresh the VDB from a Snapshot or a specific point in time. Refreshing a VDB will delete any changes that have been made to the VDB prior to the refresh operation; you are essentially resetting it to the state you select during the refresh process.
Rewind	Rewinding a VDB rolls it back to a previous point in its Timeflow. The VDB will no longer contain changes that occurred after the rewind point.
Staging Push	Allow users to control the staging database restore and the state of the restored staging database.
Disable	Remove all evidence of a specific VDB from a target environment. This action deletes the database and network pathway(s) on the target host. But it leaves the current VDB filesystem and snapshot(s) intact in Delphix Virtualization Engine so that VDB can later be enabled on a chosen target environment.
Enable	Establish the network pathway for the VDB between Delphix Virtualization Engine and the appropriate target environment, then start the specific VDB database.
Stop	Leave network pathway (e.g., NFS or iSCSI) in place on the target environment for VDB, but stop the specific VDB database on the target environment.
Start	Start a specific VDB database on the target environment. This will result in no action for the network pathway and is assumed to be in place already.

Delete	Remove all evidence of specific VDB from the target environment and the Delphix Virtualization Engine. This operation is intentionally destructive.
--------	---

Delphix Virtualization Users and Privileges

Object	User Privileges	Group Privileges
Reader	Access statistics on the dSource, VDB, or Snapshot such as usage, history, and space consumption	Access statistics on all dSources, VDBs, or Snapshots in the group such as usage, history, and space consumption
Provisioner	<ul style="list-style-type: none"> Access statistics on the dSource, VDB, or Snapshot such as usage, history, and space consumption Provision VDBs from owned dSources and VDBs 	<ul style="list-style-type: none"> Access statistics on all dSources, VDBs, or Snapshots in the group such as usage, history, and space consumption Provision VDBs from all dSources and VDBs in the group
Owner	<ul style="list-style-type: none"> Provision VDBs from owned dSources and VDBs Perform Virtual to Physical (V2P) from owned dSources Access the same statistics as a Reader Refresh or Rollback VDBs Take Snapshots of dSources and VDBs 	<ul style="list-style-type: none"> Provision VDBs from all dSources and VDBs in the group Refresh or Rollback all VDBs in the group Snapshot all dSources and VDBs in the group Perform Virtual to Physical (V2P) from owned dSources Apply custom policies to dSources and VDBs Create template policies for the group Assign Owner privileges for dSources and VDBs Access the same statistics as a Provisioner, Data Operator, or Reader
Data operator	<ul style="list-style-type: none"> Access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption Refresh or rollback VDBs 	<ul style="list-style-type: none"> Access statistics on all dSources, VDBs, or snapshots in the group such as usage, history, and space consumption Refresh or rollback all VDBs in the group

Administrator	Manage all data objects: dSources, virtual databases (VDBs), users, groups, and related policies and resources." The `admin` user manages the Delphix Virtualization Engine using either the browser-based Delphix Management application or the Command Line Interface (CLI).	Has privileges over all objects and users.
sysadmin	Can perform typical system administration duties such as: modifying NTP, SNMP, SMTP settings; managing storage; downloading support logs for the Delphix Virtualization Engine, and performing upgrades and patches. The sysadmin user launches the initial Delphix Setup configuration application and has access to the Command Line Interface (CLI).	Has privileges for storage, upgrades, network, etc.

Types of Notification

Type	Notification
Event	Completion of some action in the Delphix Virtualization Engine. Examples include user-initiated tasks such as snapshots or VDB provisioning, policy-based tasks, and background monitoring and maintenance tasks.
Alert	Caused by a single event on a Delphix Virtualization Engine. Also known as a System Event, and viewable through the System Event Viewer. Examples include warnings on source/target environment settings, recoverable errors, or incorrect connection settings. Alert Levels: Informational, Warning, Critical
Fault	A persistent event on a Delphix Virtualization Engine that remains until the issue is resolved. The fault may be marked resolved automatically or require that it be resolved manually. Selecting to Ignore a fault will also ignore future faults of that exact type against the same object. System faults describe states and configurations that may negatively impact the functionality of the Delphix Virtualization Engine and which can only be resolved through active user intervention. Examples: Delphix Virtualization Engine storage failure, communication failures between the Delphix Virtualization Engine and a source or target environment/host Fault Levels: Warning, Critical

Delphix Self-Service Terms

Term	Explanation
Administrator	<p>Has full access to all report data and can configure and administer Delphix Self-Service. Additionally, can use the Delphix Virtualization Engine to:</p> <ul style="list-style-type: none"> • add/delete reports • add/delete users • change tunable settings • add/delete tags • create and assign data templates and containers
Bookmark	<p>A logical reference to a point in time on a branch. You can use it as a point from which to fork new branches. It can also be the target of policies – for example, you can arrange to keep this bookmark for two years.</p> <p>Bookmarks are a way to mark and name a particular moment of data on a timeline. You can restore the active branch's timeline to the moment of data marked with a bookmark. You can also share bookmarks with other Delphix Self-Service users, which allows them to restore their own active branches to the moment of data in your container. The data represented by a bookmark is protected and will not be deleted until the bookmark is deleted.</p>
Branches	<p>Branches are task-specific groupings you can create within a data container. A branch is used to track a logical task and contains a timeline of the historical data for that task. As you work within your data container, you can create more branches overtime to run or complete separate tasks.</p> <p>Branches represent a logical sequence of activity, separate from the underlying data lineage. This is the main concept introduced in the core engine and forms the basis of many higher-level primitives. Branches:</p> <ul style="list-style-type: none"> • Can have only one timeline active at any time • Can be user-visible (e.g. exported to a user target) or implementation (e.g. just a staging source to run a series of transformations)
Branch group/target group	<p>A collection of multiple Branches that are treated as a single entity. The system can determine compatibility automatically, or a template can be used to create more complex orchestration.</p>
Branch timeline	<p>A dynamic point-in-time interface for user actions within the Branch. Common activities include re-setting Data Sources to run a test, refreshing the Data Container with the most current source data, and bookmarking data to share or track interesting moments of time along the branch timeline.</p>

Data container	<p>Consists of one or more Data Sources, such as databases, application binaries, or other application data. Allows users to:</p> <ul style="list-style-type: none"> • Undo any changes to their application data in seconds or minutes • Have immediate access to any version of their data over the course of their project • Share their data with other people on their team, without needing to relinquish control of their own container • Refresh their data from production data without waiting for an overworked DBA
Data template	<p>Created by the Delphix Administrator, data templates consist of the data sources users need in order to manage their data playground and their testing and/or development environments. Data templates serve as the parent for a set of data containers that the administrator assigns to Delphix Self-Service users. Additionally, data templates enforce the boundaries for how data is shared. Data can only be shared directly with other users whose containers were created from the same parent data template.</p>
Data User	<p>Delphix Self-Service data users have access to production data provided in a data container. The data container provides these users with a playground in which to work with data using the Self-Service Toolbar.</p>

Related Topics

- [Delphix Engine Overview](#)
- [Getting Started with Data Sources](#)

Product Icon Reference

This topic illustrates the icons that appear on dSources and virtual databases (VDBs) in the Delphix Engine graphic user interface and describes the meaning of each, along with tips for clearing those that represent errors.

Icon	Description
	Environment
	Host
	Cluster
	dSource
	VDB
	Masked VDB
	vFiles
	Live Source
	Universal control for adding, creating users and objects
	Provision a Virtual Database (VDB)
	Refresh a VDB
	Take a snapshot
	Rewind or roll back a VDB to a point in time
	Copy Virtual Database to a physical database (V2P)

Icon	Description
	Open Log Sync to create a dataset from a point in time
	Filter control to narrow the results in a list view
	Toggle to expand the current state
	Toggle to collapse current state
	Timeflow view selector
	There is a warning fault associated with the dataset
	There is a critical fault associated with the dataset

Related Topics

- [Delphix Engine Overview](#)
- [Getting Started with Data Sources](#)


Deployment

 Unknown macro: 'navbox'

Standard Deployment Architecture

Overview

At a high level, Delphix Virtualization functions to efficiently ingest data off of a source host by creating a compressed copy of that data, called a dSource, on the Delphix Engine. From there, the Delphix Engine can easily create multiple virtual databases (VDBs) on a target host with a marginal increase in storage. From an architecture perspective, the diagram below details how an engine interacts with an organizations infrastructure.



Architectural diagram of a common Delphix deployment (configurations vary depending on RDBMS platforms).

On an infrastructure level, Delphix Virtualization relies on a Source, Staging, and Target Hosts, in addition to the Delphix Engine needing specific compute requirements (outlined in the platform-specific requirements section). Storage Requirements, like compute, vary by host type and are outlined in subsequent sections. Staging and Target host storage is provided from the Delphix Engine, which is mounted over the network similar to any Target host (NFS/iSCSI).

Related Topics

- [Deployment](#)
- [Checklist of Information Required for Installation and Configuration](#)

Checklist of Information Required for Installation and Configuration

Overview

This article describes the information required for the initial installation and configuration of the Delphix Engine.

i The **InstallWorkbooks** articles list all the values that need to be captured, as well as sample values. Choose the workbook for the corresponding hypervisor. In addition to required values, the InstallWorkbooks also lists useful data to document as part of the installation.

- [InstallWorkbook AWS](#)
- [InstallWorkbook VMware](#)
- [InstallWorkbook Azure](#)

Hypervisor Specific Options

As a virtual appliance, Delphix supports a number of different hypervisors and configurations. The choice of hypervisor and network configuration will determine what configurations are needed for the first step in the installation of Delphix.

Attribute	Sample Data	Used with the following Hypervisor Options			
		VMware with Static IP	VMware with DHCP	AWS	Azure
Name	mydelphix1	*	*	*	*
Gateway IP	192.168.0.1	*	*	*	*
Root Volume Size	300G	*	*	*	*
ESX Hostname	thx1138	*	*		
Domain	mydomain.com	*		*	*
DNS server IP(s), comma-separated	10.80.1.1, 10.80.1.2	*		*	*
Static IP and Subnet Mask in CIDR Notation	172.16.180.130/24	*			

Attribute	Sample Data	Used with the following Hypervisor Options			
Number of vSCSI Controllers	4	*			
VMDK/RDMs and Sizes	SAN 03- 3 TB, SAN04-3 TB, SAN05-3 TB & SAN06-3 TB	*			
Virtual Private Cloud (VPC) Name	vpc-673822a23			*	
Subnet	subnet-748391e26			*	
Auto-assign Public IP	disable			*	
EC2 instance type	r3.8xlarge			*	
Security Groups	sg-78sdg99879, sg-8798s77009			*	
Root Volume Type	Magnetic			*	
Number of EBS volumes	4			*	
Size of EBS volumes	500G			*	
Type of EBS volumes	Provisioned IOPS SSD			*	
IOPS set on each volume	5,000			*	
Encryption on EBS volumes	TRUE			*	

Attribute	Sample Data	Used with the following Hypervisor Options			
Virtual Network (vnet)					*
Subnet					*
Network Interface w/ Static Public IP					*
VM Size	Example: D8s_v3				*
Network Security Group					*
OS Disk	300 GB - Premium SSD				*
Data Disks	Premium SSDs				*

Optional Information Required for Initial Configuration

Class	Attribute	Sample Data
Time	NTP Servers (Highly Recommended)	172.16.180.10, 172.16.180.11
	Timezone	US/Pacific
Serviceability (Highly Recommended)		
	Enable Phone Home? (Highly Recommended)	Yes
	Web Proxy Server IP Address	192.168.1.200
	Web Proxy Server Port	8080
	Web Proxy Server Username (If Required)	user

Class	Attribute	Sample Data
	Web Proxy Server Password (If Required)	***
	Enable SMTP Server (Highly Recommended)	Yes
	SMTP Server IP Address	192.168.22.22
	SMTP Server Port Number	25
	SMTP Server From Email Address	delphix@hostname.mydomain.com
	SMTP Server User (Optional)	user
	SMTP Server Password (Optional)	***
	SMTP Use TLS Authentication? (Optional, Default=No)	No
	SMTP Send Timeout (Optional, Default=60)	60
Authentication (Highly Recommended)		
	LDAP Server IP Address	192.168.7.7
	LDAP Server Port	389 (636 for SSL)
	Protect LDAP traffic with SSL/TLS (true false)	false
	LDAP SSL Authentication (SIMPLE DIGEST_MD5)	SIMPLE
SNMP Integration		
	SNMP Server IP Address	192.168.77.156

Class	Attribute	Sample Data
	SNMP Server Port	162
	SNMP Community	mycommunity
	SNMP Use INFORM Instead of TRAP	false
	SNMP Severity Level (INFORMATIONAL, WARNING, CRITICAL)	CRITICAL
Splunk Integration (Highly Recommended)		
	Splunk Server IP address	192.168.8.8
	Splunk Server HEC Port Number	8088
	Splunk HEC Token	12345678-1234-1234-1234-123456 7890AB
	Index Name for Events	delphix_events
	Index Name for Metrics	delphix_metrics
ADMIN		
	Email Address	u12345@mydomain.com
SYSADMIN		
	Email Address	u12345@mydomain.com

Related Topics

- [Standard Deployment Architecture](#)
- [Deployment](#)

Network Connectivity Requirements

Overview

This topic covers the general network and connectivity requirements for the Delphix Engine, including connection requirements, port allocation, and firewall and Intrusion Detection System (IDS) considerations. For platform-specific network and connectivity requirements, see the relevant topics under the **Requirements** article for each platform.

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target (see Configuring Replication)
TCP	50001	Connections to source and target environments for network performance tests

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine

Protocol	Port Number	Use
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

Related Topics

- [Configuring Replication](#)
- [Deployment](#)

Setting Up Network Access to the Delphix Engine

Overview

This article outlines the procedure for setting up network access to the Delphix Engine. Follow the initial installation instructions in [Installing the Delphix Engine](#) before addressing this procedure.

Procedure

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online. This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin` for the username and the password.
5. A description of available network settings and instructions for editing will be shown.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

`defaultRoute` IP address of the gateway **for** the **default** route -- **for** example, `"1.2.3.4."`

`dhcp` Boolean value indicating whether DHCP should be used **for** the primary **interface**. Setting **this** value to `"true"` will cause all other properties (address, hostname, and DNS) to be derived from the DHCP response

`dnsDomain` DNS Domain -- **for** example, `"delphix.com"`

`dnsServers` DNS server(s) as a list of IP addresses -- **for** example, `"1.2.3.4,5.6.7.8."`

`hostname` Canonical system hostname, used in alert and other logs -- **for** example, `"myserver"`

`primaryAddress` Static address **for** the primary **interface** in CIDR notation -- **for** example, `"1.2.3.4/22"`

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24

```

6. Configure the `hostname` . If using DHCP, this step can be skipped.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same hostname entered during the server installation.

7. Configure DNS. If using DHCP, this step can be skipped.

```

delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]

```

8. Configure either a static or DHCP address:

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```

delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>

```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`).

9. Configure a default gateway. If using DHCP, this step can be skipped.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

10. Commit the changes. Use the `get` command prior to committing to verify the desired configuration.

```

delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the
browser at:

```

```
http://<address>
```

Type `'exit'` to disconnect, or any other commands to **continue** using the CLI.

11. Check that the Delphix Engine can now be accessed through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
12. Exit setup.

```
delphix> exit
```

Related Topics

- [Upgrade](#)

Installation and Initial System Configurations

This section covers the following topics:

- [Initial Setup](#)
- [Customizing the Delphix Engine System Settings](#)
- [Installing an OVA or AMI](#)

Initial Setup

Overview

Delphix runs as a virtual appliance deployed in various types of platforms, as outlined in these **Deployment** articles. When first logging into an instance, the setup wizard will help with initial configurations for network, storage, authentication, and more. This article describes each step in the setup process and will outline and the various options available.

When first connecting to the Delphix Engine via any supported browser, enter the default sysadmin login; the credentials for username/password are sysadmin/sysadmin. On the first login, there will be a prompt to change the initial default password for security purposes.

Info

A login failure issue could occur if the Delphix Engine clock is not in sync with the IDP clock. To resolve the issue, either use the NTP clock or set up the skew time property in SSO (Single Sign-On) configuration.

Welcome

The Welcome tab asks users to select the engine type for setup, choose from Virtualization or Masking. This document explains the setup for Virtualization engines. For Masking engines, please visit the [Masking documentation](#).

Virtualization Setup

- Welcome
- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

Welcome

Choose engine type to setup:

Virtualization
 Masking

This wizard will step you through the setup. During this process you will complete the following:

- Create your password for the default "sysadmin" user
- Set the system time
- Configure network and services
- Configure the storage pool
- Configure proxies, SMTP, and LDAP (these are optional)
- Register your software

After setup is complete, you will have two administrators defined:

- The system administrator, "**sysadmin**" with the password you defined. This will be the system administrator for the instance.
- The engine administrator, "**admin**" with the password you defined. This is typically a DBA who will administer all the data managed by the instance.

When setup is complete, log in as engine administrator to begin using your engine.

Administrators

Each Delphix Virtualization engine has two default accounts:

- System Administrator: 'sysadmin' with a password that users can define. This will be the system administrator for the instance.
- Engine Administrator: 'admin' with a password that users can define. This is typically a DBA who will administer all the data managed by the instance.

Provide an email address and password for both users at the Administrators step of the setup.

Each Delphix Masking engine has the following default account:

- Masking Administrator - "admin" with a user-defined password. This will be the Masking administrator responsible for setting up users and other administrative actions in Masking. This is only applicable if Masking was selected.

Time

The Delphix Engine leverages its time setting to determine policies and actions that take place within the application. Manually set the time or choose from an NTP server, an explanation of these options are shown below.

Option	Notes
Set NTP Server (recommended)	After selecting this option, select an NTP server from the list, or click Add NTP Server to manually enter one or more server(s). When configuring a Delphix Engine on VMware, be sure to configure the NTP client on the host to use the same servers that are entered here.
Manually Select Time and Date	Click Use Browser Time and Date option to set the system time, or select the date and time by using the calendar and clock icons. If this option is selected, the date and time will persist as the local time, even if time zones are changed.

Network

The initial network configuration will be pre-populated based on the deployment platform used for Delphix. For VMware deployments, Delphix defaults to the VMXNET3 network adapter.

Select 'Settings' for each Network Interface to manage the following options:

Option	Notes
DHCP or Static network addressing	For Static addressing, enter an IP Address and Subnet Mask. The static IP address must be specified in CIDR notation (for example, 192.168.1.2/24).
Jumbo Frames	This setting is highly recommended. VMXNET3 supports Ethernet jumbo frames, which can be used to maximize throughput and minimize CPU utilization.
Routing	A default gateway will be specified in this section.

DNS Services	Enter a DNS Domain Name and DNS Server to be used for this engine.
--------------	--

Network Security

Delphix installs certificates signed by the engine's Certificate Authority. Users have the ability to manage their own certificates for HTTPS and DSP (Delphix Session Protocol) connections to and from the Delphix Engine. You can add or modify certificates and certificate signing requests (CSRs) via the '...' option.

When you update the Certificate Authority certificate, your HTTPS and DSP certificates will be automatically updated.

For more information please refer to [Certificate Management](#) in the Security section of this documentation.

Storage

The Delphix Engine automatically discovers and displays storage devices. For each device, confirm that Usage Assignment is set to Data.

You can associate additional storage devices with the Delphix Engine after initial configuration, as described in [Adding and Expanding Storage Devices](#).

There are two options for storage disk usage assignment:

1. **Data:** Once you set the storage unit assignment for a disk to Data and save the configuration, you cannot change it again.
2. **Unassigned:** These are disks being held for later use.

Configure at least four disks for the storage of user data. This makes the Delphix Engine storage manager function more efficient since duplicated metadata can be distributed across multiple disks.

Outbound Connectivity

Web Proxy

If a Web Proxy Server is necessary for your environment, select 'Configure web proxy' and enter the hostname and credentials for that server.

Phone Home

The support and phone home bundles contain metadata from the Delphix Engine, but do not include the customer data that has been ingested into the engine. Redaction of known PII data (e.g. names and email addresses for Delphix users) is done on-engine before bundles are uploaded, and again after bundles are uploaded to Delphix to ensure that the latest redaction rules are applied to each bundle without requiring the engine to be continually upgraded. There may be some limited environment data in the bundle (e.g., IP addresses and database names) that are needed for debugging purposes. Support bundles are automatically deleted within 30 days after the support case is closed, or 30 days after upload, whichever comes later.

Enabling/Disabling Phonehome

Enabling this option sends information to Delphix periodically over HTTPS (SSL). This data is securely managed by the internal team for product analysis and improvements. This feature requires a connection to the internet and will use the Web Proxy Server configuration.

Perform the following steps to enable/disable phonehome.

1. Login to the Delphix Virtualization engine setup using the sysadmin credentials.
2. From the Outbound Connectivity widget, click **Modify**.
3. To enable phonehome, select the checkbox before the **Enable phone home service** option. If enabled, this service will automatically send a stream of anonymous, non-personal metadata describing user interaction with the product's user interface.
4. To disable, deselect the checkbox before the **Enable phone home service** option.
5. Click **Save** to save for the settings.

User-click Analytics

User-click Analytics is a lightweight method to capture how users interact with Delphix product UIs, allowing Delphix to collect browser-based, user-click data. Delphix does not collect, transmit, or store any personally identifiable information (PII) such as email addresses, IP addresses, usernames, etc.

SMTP Server

Select **Use an SMTP Server** and enter the server name or IP address to enable email notifications for events and alerts. When a critical fault occurs with the Delphix Engine, it will automatically send an email alert to the admin user. Make sure to configure the SMTP server so that alert emails can be sent to this user. See **System Faults** for more information.

Authentication

On the Authentication page, configure authentication protocols such as LDAP and SAML/SSO. See [User and Authentication Management](#) for further details.

LDAP

To avoid configuration issues, consult with the lightweight directory access protocol (LDAP) administrator before attempting to set up LDAP authentication of users for the Delphix Engine. When configuring LDAP, provide an LDAP Server. Two authentication methods are currently supported: **SIMPLE** and **DIGEST_MD5**.

Select to Protect LDAP traffic with SSL/TLS if desired. This option requires to import the server certificate. If LDAP has been set up as an authentication service for the Delphix Engine, add new users with LDAP as their authentication mechanism. For more information, see the [User Groups](#) article.

SAML/SSO

To enable SAML/SSO, there are two properties to set:

1. **Audience Restriction:** The audience restriction must be set to the entity id configured in the Delphix Server via the Delphix Setup. Its default value is **https://<Delphix Server ID>**, where <Delphix Server ID> is a 36-character hexadecimal string of the form xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx. See [Determining the Delphix Server ID and Host Name](#) for more on the Delphix Server ID. If the Delphix Engine does not exist or is unreachable, enter a temporary value (such as delphix-sp-id) which must later be replaced by the actual Delphix Server ID.
2. **IdP Metadata:** an XML document that must be exported from the application created in the IdP. Paste its contents into the provided field.

Kerberos

The Kerberos page allows for Kerberos authentication to communicate between hosts connected with Delphix. Enabling this option will allow Kerberos key-based authentication when adding new environments to Delphix.

1. **Realm:** the domain over which a Kerberos authentication server has the authority to authenticate a user, host, or service.
2. **Principal:** a unique identity to which Kerberos can assign tickets.

3. **Keytab:** a file containing pairs of Kerberos principals and encrypted keys (which are derived from the Kerberos password).

Registration

As described in [Registration Management](#), registration allows Delphix Support to access the engine, properly diagnose, and identify any issues during support cases. If the Delphix Engine has access to the internet, auto-register the Delphix Engine with Delphix Support credentials in the 'Online Registration' section.

If external connectivity is not immediately available, perform the manual registration.

1. Copy the Delphix Engine registration code displayed.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, e-mail the registration code to an externally accessible email account.
3. On a machine with access to the internet, use a browser to navigate to the **Delphix Registration Portal** at <http://register.delphix.com>.
4. Log in with Delphix support credentials.
5. Paste the Registration Code.
6. Click **Register**.

 The Delphix Engine will work without registration, but it is recommended to register each engine as part of the setup. Failing to register the Delphix Engine will impact its supportability.

Summary

The final summary tab will enable a review of the configurations for each page in the setup tutorial. Confirm that everything looks correct, and click submit to complete the setup.

After Setup

- After the configuration is complete, the Delphix Engine will restart and launch the browser-based Delphix Management application.
- After the Delphix Management application launches, the admin can log in using the initial default username admin and the initial default password Delphix. On the first login, there will be a prompt to change the initial password.

Related Topics

- [Customizing the Delphix Engine System Settings](#)

Customizing the Delphix Engine System Settings

Overview

This article describes how to customize the initial system setup requirements for memory, number of CPUs, storage disks, and network configuration. The OVA file used to install the Delphix Engine is configured for the minimum system requirements. These can be customized to match the capabilities of a specific system.

Procedure

1. Shut down the guest operating system and power off the Delphix Engine.
2. Under **Getting Started**, select **Edit Virtual Machine Settings**.
3. The system settings can now be customized.

Setting	Options
Memory Size	Set to 64GB or larger based on sizing analysis. In the Resource Allocation panel, ensure that Reserve all guest memory is checked. To adjust the resource allocation of a VM without rebooting it, Delphix supports "hot-plugging" CPUs and memory to ZFS. Please note that this is only available on ESX. For more information, refer to Enabling Hot-adding .
Number of CPUs	Allocate 8 vCPUs or more based on Delphix licensing. vCPUs should be fully reserved to ensure that the Delphix Engine does not compete for CPU cycles on an overcommitted host.
Supported Controllers	Delphix supports Paravirtual SCSI (PVSCSI) controllers or LSI Parallel controllers. A mix of different types of SCSI controllers is not supported within the engine. It is recommended to use PVSCSI controllers by default. Although PVSCSI supports high throughput with minimal processing cost, the performance improvements on Delphix Engine may vary from case to case.
Disks for Data Storage	Add virtual disks to provide storage for user data such as dSources and VDBs. The underlying storage must be redundant. Add a minimum of 150GB per storage disk. All virtual disks should be the same size and have the same performance characteristics. If using VMFS, use thick provisioned, lazy zeroed disks. To alleviate IO bottlenecks at the virtual controller layer, spread the virtual disks across all four virtual SCSI controllers.

Setting	Options
Data Storage Multipathing Policy	<p>For EMC storage, the multipathing policy should always be set to round-robin (default for 5.x). Additionally, change the IO Operation Limit from the default of 1000 to 1. This should be strongly considered for other storage platforms as well.</p> <p>See VMware KB article EMC VMAX and DMX Symmetrix Storage Array Recommendations for Optimal Performance on VMware ESXi/ESX.</p>
Network	<p>The network configuration is set to have a VMXNET3 network adapter. VMXNET3 is a tuned network interface that is included with the VMtools provided in the OVA file. It will be assigned to VM Network.</p> <p>JUMBO Frames VMXNET3 supports Ethernet jumbo frames, this can be used to maximize throughput and minimize CPU utilization.</p> <p>Adding Link Aggregation via VMware NIC Teaming To increase throughput or for failover, add multiple physical NICs to the vSwitch that is connected to the Delphix Engine. To increase throughput, NIC Teaming must use the Route Based on IP Hash protocol for load balancing.</p> <p>Dedicate Physical NICs to the Delphix Engine For best performance, assign the Delphix Engine to network adapters that are used exclusively by Delphix.</p>

Enabling Hot-adding

To manually enable hot-adding in the VCenter console, use the following steps.

1. Select the Delphix Engine.
2. Right-click on the engine and select **Edit Settings**.
3. Edit settings:
 - a. To enable CPU hot-add, navigate to the CPU section and select the **CPU Hot Plug** checkbox.
 - b. To enable memory hot-add, navigate to the memory section and select the **Memory Hot Plug** checkbox.
4. Modify the appropriate setting to the new value; note that only expansion of CPU and memory is supported, and attempts to reduce the values available to the Delphix Engine will fail.
5. Click **OK** to confirm the changes.

Currently, the UI does not update with new memory and CPU information; verifying that the new resources are in use requires connecting to the OS shell.

For customers migrating from versions before 6.0, an extra step is required to enable CPU/memory hot-plug.

1. After migrating, turn off the VM.
2. Open **VM Settings**.
3. Select **VM Options**.
4. Under **General**, update the guest OS to Ubuntu (64-bit).
5. Under **Advanced**, select **Edit Configuration**.
6. Set **disk.EnableUUID** to TRUE.
7. Follow normal hot-plug enable steps.

 For customers that are upgrading their existing engines from versions below 6.0.6.0, two vSphere properties need to be enabled manually after the upgrade from vSphere in order to use this feature. The properties are **cpuHotAddEnabled** and **memoryHotAddEnabled**. For migrations from Illumos to Linux, the guest OS and version in vSphere also needs to be updated to Ubuntu, in addition to the above two properties. Without this, it has been noticed that hot-adding memory is not allowed from vSphere, even if memoryHotAddEnabled is enabled.

Post-Requisites

- After making any changes to the system settings, power on the Delphix Engine again and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

Related Topics

- [Initial Setup](#)
- [Customizing the Delphix Engine System Settings](#)
- [Setting Up Network Access to the Delphix Engine](#)

Installing an OVA or AMI

Overview

This article outlines how to install the Delphix Engine using an OVA or AMI file.

Procedure to Install an OVA

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. A support login is needed from the Delphix representative or welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where the Delphix Engine will be installed.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine. This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** to be used. If using multiple physical NICs for link aggregation, vSphere NIC teaming must be used. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**. The installation will begin and the Delphix Engine will be created in the location specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

Procedure to Install an AMI

Use the Delphix-supplied AMI file to install the Delphix Engine.

1. On the Delphix download site, click the AMI to be shared and accept the Delphix License agreement. Alternatively, follow a link given by the Delphix solutions architect.
2. On the **Amazon Web Services Account Details** form presented:
 - a. Enter the **AWS Account Identifier**, which can be found here: <https://console.aws.amazon.com/billing/home?#/account>. To use the **GovCloud AWS Region**, be sure to enter the ID for the AWS Account which has GovCloud enabled.
 - b. Select which **AWS Region** the AMI should be shared in. To have the AMI shared in a different region, contact the Delphix account representative for the proper arrangements.
3. Click **Share**. The Delphix Engine will appear in the list of AMIs in AWS momentarily.

Validating Host Deployment with HostChecker

Overview

Delphix has developed a hostchecker script that contains standardized checks for source and target hosts - these checks generally fall into three buckets:

- OS and Host permissions/access
- Network Port Checks
- DB-specific functionality

OS and Host permissions/access and network port checks can (and should) be performed prior to Delphix installation to ensure a smooth deployment.

Each DB should have a specifically associated hostchecker, there is detailed documentation on the DB-specific hostchecker page.

Procedure

1. Download the appropriate **HostChecker tarball** for the engine from <https://download.delphix.com/>. Tarballs follow the naming convention "hostchecker_<OS>_<processor>.tar". For example, if validating a linux x86 host, download the hostchecker_linux_x86.tar tarball.
2. Create a working directory and extract the **HostChecker files** from the **HostChecker tarball**.

```
mkdir dlp-host-checker
cd dlp-host-checker/
tar -xf hostchecker_linux_x86.tar
```

3. Run the `sh` script contained within:

```
sh hostchecker.sh
```

This will extract the JDK included in the tarball (if necessary) and invoke the HostChecker.

```
ora10205@bbdhcp:/home/ora10205/hostchecker-> sh hostchecker.sh
Extracting the JDK from the tarball jdk-6u45-linux-i586.tar.gz.
```

Do not run the HostChecker as root; this will cause misleading or incorrect results from many of the checks.

4. Select which **checks** to run. Note that it is possible to run checks without spawning the interface. Enter `--help` to get a list of arguments that can pass to the HostChecker.
5. As the checks are made, enter the requested **arguments**.
6. Read the output of the check. The general format is that severity increases farther down the output. First comes informational output, then warnings, then errors. If you see a message that starts with `Internal Error`, forward it to Delphix Support immediately. This represents a potential bug in the HostChecker, and not necessarily a problem with your environment.
7. Error or warning messages will explain any possible problems and how to address them. Resolve the issues that the HostChecker describes. Do not be surprised or undo your work if more errors appear the next time you run HostChecker, because the error you just fixed may have been masking other problems.
8. Repeat steps 3 - 7 until all the checks return no errors or warnings.

Deployment for VMware

Overview

This article outlines the requirements for deploying the Delphix Engine on VMware (including supported versions and instance configurations), as well as recommended configuration parameters for optimal performance.

The Delphix Engine is an intensive platform, both from a network and a storage perspective. If the Delphix Engine competes with other virtual machines on the same host for resources it will result in increased latency for all operations. As such, it is crucial that the ESXi host is not over-subscribed, as this eliminates the possibility of a lack of resources for the Delphix Engine. This includes allowing a percentage of CPU resources for the hypervisor itself as it can de-schedule an entire VM if the hypervisor is needed for managing IO or compute resources.

Supported ESX Versions

Requirements	Notes
<ul style="list-style-type: none"> VMware Cloud VMware ESX/ESXi 7.0, 7.0 U1, 7.0 U2, 7.0 U3c VMware ESX/ESXi 6.7 U3 VMware ESX/ESXi 6.5 U1, 6.5 U3 VMware ESX/ESXi 6.0 	<ul style="list-style-type: none"> More recent versions of VMware are preferred, such as ESX/ESXi 6.0 - 7.0 U3c

 If a minor release version is listed as supported, then all patch releases applicable to that minor release are certified.

Virtual CPUs

Requirements	Notes
8 vCPUs	<ul style="list-style-type: none"> CPU resource shortfalls can occur both on an over-committed host as well as competition for host resources during high IO utilization. CPU reservations are strongly recommended for the Delphix VM so that Delphix is guaranteed the full complement of vCPUs even when resources are overcommitted. It is suggested to use a single core per socket unless there are specific requirements for other VMs on the same ESXi host.
Never allocate all available physical CPUs to virtual machines	<ul style="list-style-type: none"> CPU for the ESXi Server to perform hypervisor activities must be set aside before assigning vCPUs to Delphix and other VMs. We recommend that a minimum of 8-10% of the CPUs available are reserved for hypervisor operation. (e.g. 12 vCPUs on a 128 vCore system).

Memory

Requirements	Notes
128 GB vRAM (recommended) 64GB vRAM (minimum)	<ul style="list-style-type: none"> The Delphix Engine uses its memory to cache database blocks. More memory will provide better read performance. Memory reservations are required for the Delphix VM. The performance of the Delphix Engine will be significantly impacted by the over-commitment of memory resources in the ESX Server. Reservations ensure that the Delphix Engine will not be forced to swap pages during times of memory pressure on the host. A swapped page will require orders of magnitude more time to be brought back to physical memory from the ESXi swap device.
Never allocate all the available physical memory to virtual machines.	<ul style="list-style-type: none"> The default ESX minimum free memory requirement is 6% of the total RAM. When free memory falls below 6%, ESX starts swapping out the Delphix guest OS. Delphix recommends leaving about 8-10% free to avoid swapping. For example, when running on an ESX Host with 512GB of physical memory, no more than 470GB (92%) should be allocated to the Delphix VM (and all other VMs on that host).
Memory for the ESX Server to perform hypervisor activities must be set aside before assigning memory to Delphix and other VMs.	Failure to ensure sufficient memory for the host can result in a hard memory state for all VMs on the host which will result in a block for memory allocations.

Network

Requirements	Notes
The ova is pre-configured to use one virtual ethernet adapter of type VMXNET 3.	<ul style="list-style-type: none"> Jumbo frames are highly recommended to reduce CPU utilization, decrease latency and increase network throughput. (typically 10-20% throughput improvement) If additional virtual network adapters are desired, they should also be of type VMXNET 3.
A 10GbE NIC in the ESX Server is recommended.	<ul style="list-style-type: none"> For VMs having only gigabit networks, it is possible to aggregate several physical 1GbE NICs together to increase network bandwidth (but not necessarily to reduce latency). Refer to the VMware Knowledge Base article NIC Teaming in ESXi and ESX. However, it is not recommended to aggregate NICs in the Delphix Engine VM.

<p>If the network load in the ESX Server hosting the Delphix engine VM is high, dedicate one or more physical NICs to the Delphix Engine.</p>	<ul style="list-style-type: none"> • Adding NICs only works if VMs are discovered using different interfaces. The NFS/iSCSI mounts will only use the network associated with the discovery. • See General Network and Connectivity Requirements for information about specific port configurations, and Network Performance Configuration Options for information about network performance tuning
---	--

SCSI Controller

Requirements	Notes									
<p>PVSCSI (default)/ LSI Logic Parallel</p>	<p>When adding virtual disks make sure that they are evenly distributing the load across the maximum of 4 virtual SCSI controllers. Spreading the disks across available SCSI controllers evenly will ensure optimal IO performance from the disks. For example, a VM with 4 SCSI controllers and 8 virtual disks should distribute the disks across the controllers as follows:</p> <table border="1" data-bbox="576 918 1426 1700"> <tr> <td>disk0 = SCSI(0:0) - System Disk on Controller 0 Port 0 (ignore for purposes of load balancing)</td> </tr> <tr> <td>disk1 = SCSI(0:1) - Data Disk on Controller 0 Port 1</td> </tr> <tr> <td>disk2 = SCSI(1:1) - Data Disk on Controller 1 Port 1</td> </tr> <tr> <td>disk3 = SCSI(2:1) - Data Disk on Controller 2 Port 1</td> </tr> <tr> <td>disk4 = SCSI(3:1) - Data Disk on Controller 3 Port 1</td> </tr> <tr> <td>disk1 = SCSI(0:2) - Data Disk on Controller 0 Port 2</td> </tr> <tr> <td>disk2 = SCSI(1:2) - Data Disk on Controller 1 Port 2</td> </tr> <tr> <td>disk3 = SCSI(2:2) - Data Disk on Controller 2 Port 2</td> </tr> <tr> <td>disk4 = SCSI(3:2) - Data Disk on Controller 3 Port 2</td> </tr> </table> <p>Note: For load purposes, we generally focus on the DB storage and ignore the controller placement of the system disk.</p>	disk0 = SCSI(0:0) - System Disk on Controller 0 Port 0 (ignore for purposes of load balancing)	disk1 = SCSI(0:1) - Data Disk on Controller 0 Port 1	disk2 = SCSI(1:1) - Data Disk on Controller 1 Port 1	disk3 = SCSI(2:1) - Data Disk on Controller 2 Port 1	disk4 = SCSI(3:1) - Data Disk on Controller 3 Port 1	disk1 = SCSI(0:2) - Data Disk on Controller 0 Port 2	disk2 = SCSI(1:2) - Data Disk on Controller 1 Port 2	disk3 = SCSI(2:2) - Data Disk on Controller 2 Port 2	disk4 = SCSI(3:2) - Data Disk on Controller 3 Port 2
disk0 = SCSI(0:0) - System Disk on Controller 0 Port 0 (ignore for purposes of load balancing)										
disk1 = SCSI(0:1) - Data Disk on Controller 0 Port 1										
disk2 = SCSI(1:1) - Data Disk on Controller 1 Port 1										
disk3 = SCSI(2:1) - Data Disk on Controller 2 Port 1										
disk4 = SCSI(3:1) - Data Disk on Controller 3 Port 1										
disk1 = SCSI(0:2) - Data Disk on Controller 0 Port 2										
disk2 = SCSI(1:2) - Data Disk on Controller 1 Port 2										
disk3 = SCSI(2:2) - Data Disk on Controller 2 Port 2										
disk4 = SCSI(3:2) - Data Disk on Controller 3 Port 2										

General Storage

- i** VMware offers options for disk storage, which include "Independent - Persistent" and "Independent - Non-persistent". The non-persistent setting is catastrophic to the Delphix platform when the VM reboots, thus, **Independent - Persistent** is required for installation.

Requirements	Notes
Storage used for Delphix must be provisioned from storage that provides data protection.	<p>For example, using RAID levels with data protection features, or equivalent technology.</p> <p>The Delphix engine product does not protect against data loss originating at the hypervisor or SAN layers.</p> <p>For more information refer to, Optimal Storage Configuration Parameters for the Delphix Engine.</p>

Delphix Storage Options

There are three types of data that Delphix stores on disk, which is:

1. **Delphix VM Configuration Storage:** stores data related to the configuration of the Delphix VM. VM Configuration Storage includes the VMware ESX configuration data as well as log files.
2. **Delphix Engine System Disk Storage:** stores data related to the Delphix Engine system data, such as the Delphix .ova settings.
3. **Database Storage:** stores data used by Delphix objects such as dSources and virtual databases (VDBs).

Delphix VM Configuration Storage

The Delphix VM configuration should be stored in a VMFS volume (often called a "datastore").

Requirements	Notes
The VMFS volume should have enough available space to hold all ESX configuration and log files associated with the Delphix Engine.	If a memory reservation is not enabled for the Delphix Engine (in violation of memory requirements stated above), then space for a paging area equal to the Delphix Engine's VM memory must be added to the VMFS volume containing the Delphix VM configuration data.

Delphix Engine System Disk Storage

The VMFS volume must be located on shared storage in order to use vMotion and HA features.

Requirements	Notes
--------------	-------

The Delphix Engine system disk should be stored in a VMDK.	<ul style="list-style-type: none"> The VMDK for the Delphix Engine System Disk Storage is often created in the same VMFS volume as the Delphix VM definition. In that case, the datastore must have sufficient space to hold the Delphix VM Configuration, the VMDK for the system disk, and a paging area if a memory reservation was not enabled for the Delphix Engine.
The Delphix .ova file is configured for a 127GB system drive.	<ul style="list-style-type: none"> The VMFS volume where the .ova is deployed should, therefore, have at least 127GB of free space prior to deploying the .ova.

Database Storage

Shared storage is required in order to use vMotion and HA features. In addition to making sure the latest VMware patches have been applied, check with the organizations hardware vendor for updates specific to the hardware configurations. VMDKs (Virtual Machine Disks) or RDMs (Raw Device Mappings) operating in virtual compatibility mode can be used for database storage.

Requirements	Notes
A minimum of 4 VMDKs or RDMs should be allocated for database storage.	<ul style="list-style-type: none"> Allocating a minimum of 4 VMDKs or RDMs for database storage enables the Delphix File System (DxFS) to make sure that its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.
<p>If using VMDKs:</p> <ul style="list-style-type: none"> Each VMDK should be the only VMDK in its VMFS volume The VMFS volumes should be assigned to dedicated physical LUNs on redundant storage. The VMDKs should be created with the Thick Provision Lazy Zeroed option. 	<ul style="list-style-type: none"> Provisioning VMDKs from isolated VMFS volumes on dedicated physical LUNs: <ul style="list-style-type: none"> Reduces contention for the underlying physical LUNs Eliminates contention for locks on the VMFS volumes from other VMs and/or the ESX Server Console Enables higher availability of the Delphix VM by allowing vSphere to vMotion the VM to a different ESX host in the event of a failure of the Delphix ESX host
The quantity and size of VMDKs or RDMs assigned must be identical across all 4 controllers	<ul style="list-style-type: none"> If the underlying storage array allocates physical LUNs by carving them from RAID groups, the LUNs should be allocated from different RAID groups. This eliminates contention for the underlying disks in the RAID groups as the Delphix engine distributes IO across its storage devices.
The physical LUNs used for VMFS volumes and RDMs should be of the same type in terms of performance characteristics such as latency, RPMs, and RAID level.	<ul style="list-style-type: none"> The total number of disk drives that comprise the set of physical LUNs should be capable of providing the desired aggregate I/O throughput (MB/sec) and IOPS (Input/Output Operations per Second) for all virtual databases that will be hosted by the Delphix Engine.

<p>The physical LUNs used for VMFS volumes can be thin-provisioned in the storage array.</p>	<ul style="list-style-type: none"> • If the storage array allocates physical LUNs from storage pools comprising dozens of disk drives, the LUNs should be distributed evenly across the available pools.
<p>For best performance, the LUNs used for RDMs should not be thin-provisioned in the storage array but should be thick-provisioned with a size equal to the amount of storage that will be initially allocated to the Delphix Engine. The RDM can be expanded in the future when more storage is needed.</p>	<ul style="list-style-type: none"> • Using thin-provisioned LUNs in the storage array for VMFS volumes can be useful if adding storage to the Delphix engine in the future is anticipated. In this case, the LUNs should be thin-provisioned with a size larger than the amount of storage that will be initially allocated to the Delphix Engine. When appropriate, to add more storage to the Delphix engine, use vSphere to expand the size of the VMDKs. Be sure to specify that the additional storage is also thick-provisioned and eager-zeroed.

In addition to making sure the latest VMware patches have been applied, check with the organization's hardware vendor for updates specific to the hardware configurations.

Additional VMware Configuration Notes

- Running Delphix inside of vSphere is supported.
- Using vMotion on a Delphix VM is supported.
- Device passthrough is not supported.

Procedure to Install an OVA

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).

13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

Related Topics

- [Optimal Storage Configuration Parameters for the Delphix Engine](#)
- [Architecture Checklist](#)
- [General Network and Connectivity Requirements](#)
- [Deployment](#)

Deployment for KVM

Overview

This article outlines the requirements for deploying the Delphix Engine via Linux KVM, as well as recommended configuration parameters.

Contact a Delphix representative to request this capability. Delphix will assist in assuring that all KVM requirements are met to successfully run a Delphix Engine with the most appropriate configurations for the use case.

The Delphix Engine is intensive from both a network and storage perspective. If the Delphix Engine competes with other virtual machines on the same host for resources it will result in increased latency for all operations. As such, it is crucial that your KVM host is not over-subscribed, as this eliminates the possibility of a lack of resources for the Delphix Engine. This includes allowing a percentage of CPU resources for the hypervisor itself as it can de-schedule an entire VM if the hypervisor is needed for managing I/O or compute resources.

 The KVM ecosystem includes many versions/versions. Delphix is announcing general support for KVM, but will forgo any formal certification for specific versions and rather will focus on declared support for specific Linux Kernels.

Pre-Requisites

1. The KVM provider must explicitly state that the Ubuntu 18.04/20.04 kernel is supported by their variation of KVM.

Most, if not all, KVM hypervisor providers should have public facing documentation which include support matrices for guest OS compatibility. If the KVM hypervisor provider does not explicitly state support for the Ubuntu kernel versions used in the supported Linux Delphix versions (either 18.04 currently, or 20.04 in the future), Delphix cannot support the KVM provider.

As an example, RedHat Linux (a KVM hypervisor provider), provides a publicly accessible [guest operating system support matrix](#). As of December 2021, this matrix does not include support for any Ubuntu guest OS, and therefore Delphix would also not support deployment on RedHat Linux KVM. Oracle Linux (a different KVM hypervisor provider) also provides a publicly accessible [guest operating system support matrix](#), which declares support for Ubuntu 16.04, 18.04, and 20.04. In this case, since the hypervisor provider declares support for the Delphix Ubuntu kernel(s), Delphix will support deployment on Oracle Linux KVM.

2. A check of the Delphix appliance on the variation/version of KVM.

It is **required** that a Delphix representative works with the organization to check the Delphix appliance (and/or use case) on the specified KVM variation/version. This check will ensure that the software is compatible.

Virtual CPUs

Requirements	Notes
8 vCPUs	<ul style="list-style-type: none"> • CPU resource shortfalls can occur on an over-committed host as well as competition for host resources during high IO utilization. • CPU reservations are strongly recommended for the Delphix VM so that Delphix is guaranteed the full complement of CPUs even when resources are overcommitted. • It is suggested to use a single core per socket unless there are specific requirements for other VMs on the same KVM host.
Never allocate all available physical CPUs to virtual machines.	<ul style="list-style-type: none"> • A CPU for the KVM Server to perform hypervisor activities must be set aside before assigning vCPUs to Delphix and other VMs. • We recommend that a minimum of 8-10% of the CPUs available are reserved for hypervisor operation. (e.g. 12 vCPUs on a 128 vCore system).

Memory

Requirements	Notes
128 GB vRAM (recommended) 64GB vRAM (minimum)	<ul style="list-style-type: none"> • The Delphix Engine uses its memory to cache database blocks. More memory will provide better read performance. • Memory reservations are required for the Delphix VM. The performance of the Delphix Engine will be significantly impacted by the over-commitment of memory resources in the KVM Server. • Reservations ensure that the Delphix Engine will not be forced to swap pages during times of memory pressure on the host. A swapped page will require orders of magnitude more time to be brought back to physical memory from the KVM swap device.
Memory for the KVM Server to perform hypervisor activities must be set aside before assigning memory to Delphix and other VMs.	Failure to ensure sufficient memory for the host can result in a hard memory state for all VMs on the host which will result in a block for memory allocations.

Network

Requirements	Notes
Virtual ethernet adapter requirements.	<ul style="list-style-type: none"> • Jumbo frames are highly recommended to reduce CPU utilization, decrease latency, and increase network throughput (typically 10-20% throughput improvement). • A 10GbE NIC/network is recommended for performant dSource and VDB operations.
If the network load in the KVM Server hosting the Delphix engine VM is high, dedicate one or more physical NICs to the Delphix Engine.	<ul style="list-style-type: none"> • Adding NICs only works if VMs are discovered using different interfaces. The NFS/iSCSI mounts will only use the network associated with the discovery. • See General Network and Connectivity Requirements for information about specific port configurations, and Network Performance Configuration Options for information about network performance tuning.

SCSI Controller

Notes

When adding virtual disks, make sure that they are evenly distributing the load across the maximum of 4 virtual SCSI controllers. Spreading the disks across available SCSI controllers evenly will ensure optimal IO performance from the disks. For example, a VM with 4 SCSI controllers and 8 virtual disks should distribute the disks across the controllers as follows:

disk0 = SCSI(0:0) - System Disk on Controller 0 Port 0 (ignore for purposes of load balancing)
disk1 = SCSI(0:1) - Data Disk on Controller 0 Port 1
disk2 = SCSI(1:1) - Data Disk on Controller 1 Port 1
disk3 = SCSI(2:1) - Data Disk on Controller 2 Port 1
disk4 = SCSI(3:1) - Data Disk on Controller 3 Port 1
disk5 = SCSI(0:2) - Data Disk on Controller 0 Port 2
disk6 = SCSI(1:2) - Data Disk on Controller 1 Port 2
disk7 = SCSI(2:2) - Data Disk on Controller 2 Port 2
disk8 = SCSI(3:2) - Data Disk on Controller 3 Port 2

Note: For load purposes, we generally focus on the DB storage and ignore the controller placement of the system disk.

General Storage

Delphix recommends using a minimum of four disks to run your Delphix Engine. One disk is used for the Delphix File System (DxFS) to ensure that its file systems are always consistent on disk without additional serialization. The other three or more equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queuing more I/O operations to its storage.

Requirements	Notes
Storage used for Delphix must be provisioned from storage that provides data protection.	<p>For example, using RAID levels with data protection features, or equivalent technology.</p> <p>The Delphix Engine does not protect against data loss originating at the hypervisor or SAN layers.</p> <p>For more information refer to, Optimal Storage Configuration Parameters for the Delphix Engine.</p>

Deployment for Hyper-V

Overview

This article outlines the requirements for deploying the Delphix Engine on Hyper-V (including supported versions and instance configurations), as well as recommended configuration parameters for optimal performance.

Contact a Delphix representative to request this capability. Delphix will assist in assuring that all Hyper-V requirements are met to successfully run a Delphix Engine with the most appropriate configuration for the use case.

The Delphix Engine is intensive both from a network and a storage perspective. If the Delphix Engine competes with other virtual machines on the same host for resources it will result in increased latency for all operations. As such, it is crucial that your Hyper-V host is not over-subscribed, as this eliminates the possibility of a lack of resources for the Delphix Engine. This includes allowing a percentage of CPU resources for the hypervisor itself as it can de-schedule an entire VM if the hypervisor is needed for managing I/O or compute resources.

Supported Versions

- Hyper-V Version: 10.0 and later
- Virtual Machine: Gen 1 only is supported

Virtual CPUs

Requirements	Notes
8 vCPUs	<ul style="list-style-type: none"> • CPU resource shortfalls can occur on an Over-committed host as well as competition for host resources during high IO utilization. • CPU reservations are strongly recommended for the Delphix VM so that Delphix is guaranteed the full complement of CPUs even when resources are overcommitted. • It is suggested to use a single core per socket unless there are specific requirements for other VMs on the same Hyper-V host.
Never allocate all available physical CPUs to virtual machines.	<ul style="list-style-type: none"> • CPU for the Hyper-V Server to perform hypervisor activities must be set aside before assigning vCPUs to Delphix and other VMs • We recommend that a minimum of 8-10% of the CPUs available are reserved for hypervisor operation. (e.g. 12 vCPUs on a 128 vCore system).

Memory

Requirements	Notes
<p>128 GB vRAM (recommended)</p> <p>64GB vRAM (minimum)</p>	<ul style="list-style-type: none"> The Delphix Engine uses its memory to cache database blocks. More memory will provide better read performance. Memory reservations are required for the Delphix VM. The performance of the Delphix Engine will be significantly impacted by the over-commitment of memory resources in the Hyper-V Server. Reservations ensure that the Delphix Engine will not be forced to swap pages during times of memory pressure on the host. A swapped page will require orders of magnitude more time to be brought back to physical memory from the Hyper-V swap device.
<p>Memory for the Hyper-V Server to perform hypervisor activities must be set aside before assigning memory to Delphix and other VMs.</p>	<p>Failure to ensure sufficient memory for the host can result in a hard memory state for all VMs on the host which will result in a block for memory allocations.</p>

Network

Requirements	Notes
<p>Virtual ethernet adapter requirements.</p>	<ul style="list-style-type: none"> SR-IOV is recommended for all virtual ethernet adapters that will be used for Delphix data I/O. Jumbo frames are recommended. A 10GbE NIC in the Hyper-V Server is recommended.
<p>If the network load in the Hyper-V Server hosting the Delphix engine VM is high, dedicate one or more physical NICs to the Delphix Engine.</p>	<ul style="list-style-type: none"> Adding NICs only works if VMs are discovered using different interfaces. The NFS/iSCSI mounts will only use the network associated with the discovery. See General Network and Connectivity Requirements for information about specific port configurations, and Network Performance Configuration Options for information about network performance tuning.

SCSI Controller

Requirements	Notes
IDE Controller for the boot drive SCSI Controller for database storage	<p>Per Hyper-V Storage I/O Performance Tuning Guidelines; For optimal performance, it is recommended that you attach multiple disks to a single virtual SCSI controller and create additional controllers only as they are required to scale the number of disks connected to the virtual machine.</p> <p>For example, a VM with 3 virtual disks should distribute the disks across the single SCSI controller as follows:</p> <ul style="list-style-type: none"> • IDE Controller 1 <ul style="list-style-type: none"> • Boot Drive • SCSI Controller <ul style="list-style-type: none"> • disk1 • disk2 • disk3 <p>Note: For load purposes, we generally focus on the DB storage and ignore the controller placement of the system disk.</p>

General Storage

Delphix recommends using a minimum of four disks to run your Delphix Engine. One disk is used for the Delphix File System (DxFS) to ensure that its file systems are always consistent on disk without additional serialization. The other three or more equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.

Requirements	Notes
Storage used for Delphix must be provisioned from storage that provides data protection.	<p>For example, using RAID levels with data protection features, or equivalent technology.</p> <p>The Delphix Engine does not protect against data loss originating at the hypervisor or SAN layers.</p> <p>For more information refer to, Optimal Storage Configuration Parameters for the Delphix Engine.</p>

Delphix Storage Options

There are two types of data that Delphix stores on disk which must be stored on NTFS volumes: :

1. Delphix Engine System Disk Storage: This is the system boot drive.
2. Database Storage: stores data used by Delphix objects such as dSources and virtual databases (VDBs).

Delphix Engine System Disk Storage

Requirements	Notes
The Delphix Engine disks must be stored on NTFS volume(s).	The volume for the Delphix Engine System Disk Storage is often created on the same volume as the Delphix VM definition. In that case, the volume must have sufficient space to hold the Delphix VM Configuration, the virtual disk for the system disk, and a paging area if a memory reservation was not enabled for the Delphix Engine.
The Delphix .vhdx file is configured for a 128GB system drive.	The volume where the .vhdx is deployed should, therefore, have at least 127GB of free space prior to deploying the .vhdx.

Database Storage

In addition to making sure the latest Hyper-V patches have been applied, check with your hardware vendor for updates specific to your hardware configuration. VHDXs (virtual machine disks).

Requirements	Notes
A minimum of 3 VHDXs should be allocated for database storage.	Allocating a minimum of 3 VHDXs for database storage enables the Delphix File System (DxFS) to make sure that its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.
If using VHDXs: <ul style="list-style-type: none"> Each VHDX should be the only VHDX on its NTFS volume The VHDX volumes should be assigned to dedicated physical LUNs on redundant storage. 	Provisioning VHDXs from isolated volumes on dedicated physical LUNs: <ul style="list-style-type: none"> Reduces contention for the underlying physical LUNs Eliminates contention for locks on the volumes from other VMs and/or the Hyper-V Server Console
	If the underlying storage array allocates physical LUNs by carving them from RAID groups, the LUNs should be allocated from different RAID groups. This eliminates contention for the underlying disks in the RAID groups as the Delphix engine distributes IO across its storage devices.
The physical LUNs used for NTFS volumes should be of the same type in terms of performance characteristics such as latency, RPMs, and RAID levels.	The total number of disk drives that comprise the set of physical LUNs should be capable of providing the desired aggregate I/O throughput (MB/sec) and IOPs (Input/Output Operations per Second) for all virtual databases that will be hosted by the Delphix Engine

Requirements	Notes
Disk types supported by the physical LUNs used for NTFS volumes are: Fixed-size and Dynamically expanding.	If the storage array allocates physical LUNs from storage pools comprising dozens of disk drives, the LUNs should be distributed evenly across the available pools.

Related Topics

- [Optimal Storage Configuration Parameters for the Delphix Engine](#)
- [Architecture Checklist](#)
- [General Network and Connectivity Requirements](#)
- [Procedure for Deploying with Hyper-V](#)

Procedure for Deploying with Hyper-V

Overview

This article outlines the procedure for deploying the Delphix Engine in Hyper-V.

Creating a Delphix Engine VM

1. Download the image from <https://download.delphix.com> and copy it to a VM directory.
 2. Start the Hyper-V Manager and specify **Name** and **Location**, then select **Next**.
 3. Select **Generation 1**, configure memory and then select **Next**.
Memory: 64 GB (minimum), 128 GB (recommended)
 4. Set up Networking by selecting **vNIC**, then select **Next**.
 5. Attach the downloaded image as a boot disk. Create a unique boot disk for each image. Please note that boot disks cannot be shared.
1. Use an existing virtual hard disk.
 2. Browse to the location of VM.
 3. Select the Image.
- Select **Finish**, the VM will appear in the inventory.

Customize the VM by selecting **Settings**

1. Delphix recommends having the IDE as the first device to boot from (under BIOS setting).
2. Adjust the number of CPUs (min 8).
3. Add a minimum of four equal-sized **Hard Drives** for the database storage. Use VHDX formatted disks. Recommend Fixed Size disk type. Please note that Differencing Disk Types are not supported.
4. Connect to the console and start the VM.
5. Once the installation has been completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

Related Topics

- [Deployment for Hyper-V](#)

Deployment for AWS EC2

Overview

This article outlines the virtual machine requirements, including memory and data storage, for deploying the Delphix Engine on Amazon EC2 (Elastic Cloud Compute). Once the requirements listed on this page are reviewed, refer to the next articles in the AWS EC2 Deployment topic:

- [Prerequisites to Deploying in AWS](#)
- [Procedure for Deploying in AWS](#)

Instance Types

The following is a list of instance types that are supported to deploy Delphix on AWS EC2. Delphix periodically certifies new instance types, which will be added to the list here.

Requirements	Notes
Large Memory Instances (Preferred): <ul style="list-style-type: none"> • R4 • R5n • R5b 	<ul style="list-style-type: none"> • The Delphix Engine most closely resembles a storage appliance and performs best when provisioned using a storage-optimized instance type. • Larger instance types provide more CPU, which can prevent resource shortfalls under high I/O throughput conditions. • Larger instances also provide more memory, which the Delphix Engine uses to cache database blocks. More memory will provide better read performance. • Recommended minimum size of 8 vCPU and 64G memory.

Network Configuration

It is recommended to enable Enhanced Networking to maximize performance, for more information view [Enhanced Networking on Linux](#).

Requirements	Notes
Virtual Private Cloud	The Delphix Engine and all of the source and target environments must be deployed in a VPC network to ensure that private IP addresses are static and do not change when restarting instances. When adding environments to the Delphix Engine, use the host's VPC (static private) IP addresses.
Static Public IP	The EC2 Delphix instance must be launched with a static IP address; however, the default behavior for VPC instances is to launch with a dynamic public IP address – which can change whenever the instance restarts. If using a public IP address for the Delphix Engine, static IP addresses can only be achieved by using assigned AWS Elastic IP Addresses .

Security Group Configuration	<p>The default security group will only open port 22 for SSH access. The security group must be modified to allow access to all of the networking ports used by the Delphix Engine and the various source and target engines.</p> <p>See Network Performance Configuration Options for information about network performance tuning.</p> <p>See General Network and Connectivity Requirements for information about specific port configurations.</p>
------------------------------	---

EBS Configuration

Deploying Delphix on AWS EC2 requires EBS volumes. Since EBS volumes are connected to EC2 instances via the network, other network activity on the instance can affect throughput to EBS volumes. Delphix recommends EBS General Purpose SSD (GP2 and GP3) or Provisioned IOPS SSD volumes to provide consistent and predictable storage performance. For more information on EBS volumes, see external documentation on [EBS Volume Types](#).

Recommendations	Notes
EBS Provisioned IOPS SSD Volumes (Highly Recommended)	<p>Delphix does not support the use of instance store volumes.</p> <p>Use EBS volumes with provisioned IOPs in order to provide consistent and predictable performance. The number of provisioned IOPs depends on the estimated IO workload on the Delphix Engine.</p> <p>Provisioned IOPs volumes must be configured with a volume size to provisioned IOPs per the EBS Volume Types guidelines.</p>
EBS General Purpose SSD Volumes	<p>Delphix supports GP2 and GP3 General Purpose SSDs. You can now migrate your Amazon EBS volumes from GP2 to GP3. For more details on migration, see the external documentation on EBS Volumes Migration.</p>

General Storage Configuration

Requirements	Notes
--------------	-------

General Storage Configuration	<p>Allocate initial storage equal to the size of the physical source database storage. For high redo rates and/or high DB change rates, allocate an additional 10-20 %.</p> <p>Add storage when storage capacity approaches 30% free.</p> <p>All data storage volumes must be EBS volumes.</p> <p>Delphix recommends using a minimum of four equally sized storage disks to run the Delphix Engine. This allows the Delphix File System (DxFS) to assure its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.</p>
-------------------------------	--

Additional AWS Configuration Notes

- Using storage other than EBS is not supported.
- Limits on the number of volumes are dictated by the EBS instance type, and is generally advised that over 40 can be expected to cause issue on Linux VMs. More information can be found in the [AWS Volume Limits](#) and [AWS Volume Constraints](#) articles. The maximum device limit imposed by AWS can be handled by the Delphix Engine.
- Cold HDD (sc1) volumes (not supported due to poor performance)
- Using fast storage for EBS volumes is supported and recommended, including (in order of decreasing speed):
 - Provisioned IOPs SSD volumes (recommended)
 - General Purpose SSD volumes (supported)
 - Throughput Optimized HDD (supported, not recommended due to performance)
- Use of EBS volumes encrypted at creation is supported (during initial deployment from AMI, as well as storage devices added post-deployment), but can have negative performance consequences. Conversion of existing EBS volumes is possible in AWS but is not supported in Delphix at this time.

Related Topics

- [General Network and Connectivity Requirements](#)
- [Network Performance Configuration Options](#)
- [Configuration](#)

Prerequisites to Deploying in AWS

Overview

This article outlines the prerequisites for deployment of the Delphix Engine on AWS. The setup user should have experience launching and configuring instances in the Amazon Web Services environment. Review and complete the tasks in the next section before deployment.

Prerequisites

1. Review [Checklist of Information Required for Installation and Configuration](#).
2. Make sure that the Amazon account being used to deploy the Delphix Engine has an appropriate level of enablement to subscribe to the Delphix Engine for AWS subscription.
3. Determine which virtual private cloud (VPC) is being used when launching the virtualization instance. To maximize performance, deploy the Delphix Engine instance in the same VPC/subnet in which the virtual databases (VDBs) will be created.

 Provisioning a VDB requires a compute instance running the same database engine as the source. Please note, however, that the target instance only needs storage to accommodate the OS, database platform binaries, etc., because Delphix delivers all of the data files.

4. Make sure that the necessary ports are open.

 Using the Delphix Engine for AWS will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform:

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network Access Requirements for SQL Server Environments](#)
- [Network and Connectivity Requirements for SAP ASE Base Environments](#)
- [Network and Connectivity Requirements for Db2 Environments](#)

5. Update Security Group settings to accommodate the necessary connections.
 - a. Select the same Security Group that the current (or future) non-production EC2 compute nodes utilize.
 - b. Modify the Security Group to allow access to all of the networking ports used by the Delphix Engine and the various source and target platforms. See links above for information about specific port configurations.
6. Allocate storage.
 - a. To properly size the initial storage capacity and determine the number and size of EBS Provisioned IOPs Volumes required, download and utilize the [Delphix-Dynamic-Data-Platform-Storage-Calculator](#).

 It is helpful to first create a list of the data sources intended for making dSources. A data source is typically a production database linked to the Virtualization Engine, enabling to create virtual, full, read-write copies of the source within minutes. The list should include the database name, platform (for example, Oracle or SQL Server), current size (in GB), the estimated number of virtual copies, and retention period (in days) of snapshots (backup copies).

- b. All data storage volumes must be EBS volumes. Delphix recommends using a minimum total of four disks to run the Delphix Engine. One disk is used for the boot device. The other four equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.
 - c. Provisioned IOPS EBS volumes are highly recommended.
7. During the Manual Deployment option, use the guidelines outlined in [Virtual Machine Requirements for AWS EC2 Platform](#).

Geographic Distribution in Regions and Availability Zones

The latency will be directly related to not just the Availability Zone configuration, but more specifically the geographies of those zones. The latencies can vary from tens to hundreds of milliseconds if the zones are geographically diverse (US West to US East would certainly be expected to perform better than US West to Europe). AWS advertises that all AZs in a given region are interconnected with high-bandwidth and low-latency networking per [this AWS article](#).

Applications that are performance-sensitive will benefit from colocating the target servers and Engine in the same region if possible. Another possibility is building a failover strategy for those highly sensitive servers, enabling them to failover to another AZ in the instance where an Engine needs to be failed over. The architecture selected and geographies will also naturally be dependent on your redundancy requirements (your organization may require geographically diverse failover options beyond the ~60 miles advertised within a region).

Related Topics

- [Deployment for AWS EC2](#)
- [Procedure for Deploying in AWS](#)
- [Checklist of Information Required for Installation and Configuration](#)

Procedure for Deploying in AWS

Overview

This article outlines the procedure for deploying the Delphix Engine on AWS, using the Delphix Marketplace Image or Amazon Machine Image (AMI).

Procedure to Install an AMI

Use the Delphix-supplied AMI file to install the Delphix Engine.

1. On the Delphix download site, select the AMI to be shared and accept the Delphix License agreement. Alternatively, use a link provided by the organization's Delphix solutions architect.
2. On the **Amazon Web Services Account Details** form presented:
 - a. Enter the AWS Account Identifier, which can be found here: <https://console.aws.amazon.com/billing/home?#/account>. To use the **GovCloud AWS Region**, be sure to enter the ID for the AWS Account which has GovCloud enabled.
 - b. Select which **AWS Region** that the AMI will be shared in. For an AMI shared in a different region, contact the organization's Delphix account representative to make the proper arrangements.
3. Click **Share**. The Delphix Engine will appear in the list of AMIs in AWS momentarily.

Subscribe to the Delphix Virtualization Engine Marketplace Image

1. Login to the AWS Console at <https://aws.amazon.com/console>.
2. Navigate to the AWS Marketplace.
3. In the **Search** field, enter **Delphix**.
4. Select the **Delphix Dynamic Data Platform for AWS**. Select the appropriate offer based on the data you will ingest into the Delphix platform.
5. Review the information on the initial Marketplace page.
6. Click **Continue to Subscribe**.
7. Review the software subscription information and select the **Manual Launch** tab.
8. Select **Accept Software Terms**.
Your subscription will take a few minutes to become enabled.
9. From the **Subscription** page, click **Return to Product Page**. There you will see **Launch with EC2 Console** enabled in all of the available regions.

Launching the Delphix Engine

1. In the desired region, select **Launch with EC2 Console**. Launching will redirect to the standard EC2 instance launch process.
2. Select the size of the virtual machine to deploy. For supported instance types and capabilities, see [Virtual Machine Requirements for AWS Platform](#).
3. Click **Next: Configure Instance Details**.
4. Configure instance details as follows:
 - a. Set the **number of instances** to **1**.
 - b. **Purchasing option** – Leave **Request Spot Instances** unchecked.
 - c. **Network** – Select the VPC into which the instance will be deployed.
 - d. **Subnet** – Select the Subnet into which the instance will be deployed.
 - e. **Auto-assign Public IP** – Select **Disable**.
 - f. **Placement group** – Optional. The default is **No placement group**.
 - g. **IAM Role** – None

- h. **Shutdown behavior** – Stop
- i. **Enable termination protection** – Yes (select checkbox). This ensures that the Delphix Engine is not accidentally terminated.
- j. **Monitoring** – Optional
- k. **EBS-optimized instance** – Yes (select checkbox)
- l. **Tenancy** – Shared – Run a shared hardware instance
- m. **Network Interfaces** – Eth0

 This option will not appear if a default subnet is chosen. In this case, the first option below will apply.

- i. Auto-assign (default if using DHCP)
- ii. Enter static IP in Primary IP field (recommended for consistency)

5. Click **Next: Add Storage**.
6. By default, the Delphix Virtualization instance includes a 150GB OS volume. Ensure that this volume is set as a **General Purpose SSD (GP2)** type.
7. Using the [Delphix Virtualization Engine Storage Calculator for AWS](#), enter information from the predetermined inventory into the worksheet. This will automatically calculate the number, size, and IOPS required for the EBS volumes.

Delphix Virtualization Engine - Storage Calculator for AWS Marketplace					
DB Name	DB Type	DB Source Size (GB)	Est. # of Virtual DBs	Retention for Snapshots (Days)	Delphix Storage (GB)
oradb1	Oracle	2500	5	14	3000
sqldb1	MS-SQL	1500	5	14	1800
sqldb2	MS-SQL	150	5	14	180
					0
					0
Totals		4150			4980
# of EBS Volumes					4
Size per Volume					1250
IOPS per Volume					1000

In the above example, there are three (3) source databases listed with their current sizes. Each is estimated to have 5 virtual copies as well as a 14-day retention period for snapshots/backups. The calculator indicates to add four (4) EBS-provisioned IOPS volumes, each being 1250GB with 1000 IOPS.

 This calculator is only for estimating the anticipated storage size and IOPS needs. A manual adjustment may be required prior to or after launch depending on storage and performance needs.

8. Click **Add New Volume**. The new volume will be created using the specifications in the estimation provided. The below values are sample values set for a Volume Type. See [EBS Configuration](#) for related information.
 - a. **Volume Type** – EBS
 - b. **Device** – Use default provided
 - c. **Size (GiB)** – 1250
 - d. **Volume Type** – Provisioned IOPS SSD (IO1)
Delphix deploys as an EC2 instance on AWS using EBS volumes for storage. AWS provides the following options for Elastic Block Storage (EBS) storage volumes, namely
 - i. Standard
 - ii. General Purpose SSD (GP2)
 - iii. General Purpose SSD (GP3)
 - iv. Provisioned IOPS SSD (IO1)
 - e. **IOPS** – 1000
 - f. **Delete on Termination** – Optional. If this option is not set and the Delphix Engine is terminated, the EBS volumes will persist, and they will need to be removed manually.

g. **Encrypted** – Optional

- Follow the above procedure for adding as many EBS volumes as specified in the calculator. Based on the example above, the storage configuration would look as follows:

Step 4: Add Storage
Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. [Learn more](#) about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination	Encrypted
Root	/dev/xvda	snap-1d44a945	150	General Purpose SSD (GP2)	450 / 3000	N/A	<input checked="" type="checkbox"/>	Not Encrypted
EBS	/dev/sdb	Search (case-insensit)	1250	Provisioned IOPS SSD (IO1)	1000	N/A	<input type="checkbox"/>	<input type="checkbox"/>
EBS	/dev/sdc	Search (case-insensit)	1250	Provisioned IOPS SSD (IO1)	1000	N/A	<input type="checkbox"/>	<input type="checkbox"/>
EBS	/dev/sdd	Search (case-insensit)	1250	Provisioned IOPS SSD (IO1)	1000	N/A	<input type="checkbox"/>	<input type="checkbox"/>
EBS	/dev/sde	Search (case-insensit)	1250	Provisioned IOPS SSD (IO1)	1000	N/A	<input type="checkbox"/>	<input type="checkbox"/>

Add New Volume

General Purpose (SSD) volumes provide the ability to burst to 3000 IOPS per volume, independent of volume size, to meet the performance needs of most applications and also deliver a consistent baseline of 3 IOPS/GiB. [Set my root volume to General Purpose \(SSD\).](#)

Free tier eligible customers can get up to 30 GB of EBS General Purpose (SSD) or Magnetic storage. [Learn more](#) about free usage tier eligibility and usage restrictions.

Cancel Previous **Review and Launch** Next: Tag Instance

- Verify the storage configuration matches the output from the calculator.
- Click **Next: Tag Instance**.
- To provide a name for the Delphix Virtualization Instance, enter it in the **Value** field corresponding to the **Key Name**.
- Click **Next: Configure Security Group**.
- Assign a security group. Do one of the following:
 - Create a new security group.
 - If an existing security group meets Delphix requirements, it can be selected. For more details on required and recommended ports for general Delphix usage, see [General Network and Connectivity Requirements](#).
- Click **Review and Launch**.
- Verify the Delphix Virtualization Instance details and click **Launch**.
- Proceed without a key pair.
- Select the **I acknowledge** checkbox.
- Click **Launch Instances**.
- Click **View Instances**.
- To complete the initial Delphix Virtualization Engine configuration, wait for the instance to be up and running. The status is available on the **Instances** page. Note that it can take up to 15 minutes for the first launch of Delphix, and even if the status of the server is “running” it may not yet be ready.

Configuring the Delphix Engine

- Connect to the running Delphix instance with a web browser. Use the IP address or DNS name noted in the Instance Description.
- Upon successful connection, the browser will automatically redirect to the **Delphix Setup Page**.
- Refer to the standard [product deployment instructions](#) to complete the Delphix deployment.

Logging in for the first time

On the first time login to the Delphix Engine, follow these steps:

1. Enter the default Administration User: **sysadmin**.
2. Enter the default password: the **<Instance ID>** for the Delphix Engine.

 Find the *<Instance ID>* in the **Instances** section of the AWS EC2 Management Console.

3. Once redirected, enter the following as prompted:
 - a. **Email address** for the admin account.
 - b. A new **password**. This password will be used for the Engine Administrator with the username "admin".
 - c. Confirm the new password.

 From now on, this is the password used to log in to the Delphix Engine Administrator interface to manage Datasets.

4. Click **Continue**.

Next Steps

Congratulations! The Delphix Virtualization Engine should be successfully deployed in AWS.

Use Delphix documentation to learn how to:

- configure a database source
- configure target environments
- create virtual databases (VDBs)

Related Topics

- [Deployment for AWS EC2](#)
- [Prerequisites to Deploying in AWS](#)

AWS RDS Custom for Oracle and SQL Server

Overview

The Amazon Relational Database Service (Amazon RDS) makes it easy to set up, operate, and scale a relational database in the cloud. It allows organizations to automate time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. It frees focus that can be utilized on applications, eliminating extra overhead. Amazon RDS is available for Oracle and SQL Server, amongst other database services.

Delphix AWS RDS Custom Support for Oracle and SQL Server

AWS has announced a new version of their RDS databases, thus, Delphix is adding AWS RDS Custom support for Oracle. AWS RDS Custom is an RDS service that allows Delphix to access the operating system and database level functions required to provision vDBs into the AWS RDS PaaS (Product as a Service) database service.

With AWS RDS Custom, organizations using Delphix can obtain the full benefits of PaaS and Delphix automation for their target environments.

- Users ingest data from production in RDS Oracle/SQL Server, mask the data, and provision secure copies back to RDS. All Delphix automation tools work and also manual operations via the UI.
- Delphix data compression and block sharing capabilities will allow organizations to potentially reduce their AWS RDS storage consumption.

For additional questions or use case evaluations, please contact Delphix Support.

Deployment for Microsoft Azure

Overview

This article outlines the virtual machine requirements, including memory and data storage, for deploying the Delphix Engine on the Azure Public Cloud and Government Cloud. Once the requirements listed on this page are reviewed, refer to the next articles on Azure deployment:

- [Prerequisites to Deploying in Microsoft Azure](#)
- [Procedure for Deploying in Microsoft Azure](#)

Instance Types

The following is a list of instance types that are supported to deploy Delphix on Microsoft Azure. Delphix periodically certifies new instance types, which will be added to the list here.

Requirements	Notes
General Purpose	
<ul style="list-style-type: none"> • Dsv2-series • Dsv3-series 	<p>Network bandwidth and IOPS limits are specific to each instance type:</p> <ul style="list-style-type: none"> • See DSv3 specifications for more details. <p>Recommended minimum size of 8 vCPU and 64G memory.</p>
Memory-Optimized	
<ul style="list-style-type: none"> • Esv3-series • Esv4-series • Esv5-series 	<p>Network bandwidth and IOPS limits are specific to each instance type:</p> <ul style="list-style-type: none"> • See DSv2 specifications for more details. • See ESv3 specifications for more details. <p>Recommended minimum size of 8 vCPU and 64G memory.</p>

Network Configuration

The use of Azure Accelerated Networking is supported for Delphix Engines and is recommended to maximize network performance. Accelerated Networking provides more bandwidth with more consistent and lower network latencies.

Requirements	Notes
--------------	-------

Azure Virtual Network (VNet)	The Delphix Engine must have ≤ 1 ms of latency to the Target and Staging environments. Longer latencies are permitted between the Delphix Engine and Source environments, but may impact Snapsync timings. VNET peering may be used, as long as latency requirements are met. For very low latency requirements, consider Proximity Placement Groups for Delphix Engine, Target, and Staging.
Network Security Group (NSG)	The security group must be modified that allows access to all networking ports used by the Delphix Engine and the various source/target platforms. See General Network and Connectivity Requirements for information about specific port configurations. See Network Performance Configuration Options for information about network performance tuning.

Storage Configuration

Requirements	Notes
General Storage Configuration	<ul style="list-style-type: none"> • Allocate initial storage equal to the size of the physical source databases. For high redo rates and/or high DB change rates, allocate an additional 10-20% storage. • Add storage when storage capacity approaches 30% free. • Delphix recommends using a minimum of four equally sized storage disks to run the Delphix Engine. This allows the Delphix File System (DxFS) to assure its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage. • Maximize Delphix Engine RAM for a larger system cache to service reads. • See Delphix Storage Migration.
Azure Premium Storage	<ul style="list-style-type: none"> • Premium storage utilizes solid-state drives (SSDs). Ultra disk is also supported. • Devices up to 32TB are supported. • A maximum of 256TB is supported. • Cache setting of NONE. • I/O requests of up to 256 kilobytes (KB) are counted as a single I/O operation (IOP) for provisioned IOPS volumes. • IOPS vary based on storage size with a maximum of 20,000 IOPS.

Extensions

Azure VM Extensions are not currently supported.

Related Topics

- [Prerequisites to Deploying in Azure](#)
- [Procedure for Deploying in Azure](#)
- [General Network and Connectivity Requirements](#)
- [Network Performance Configuration Options](#)
- [Delphix Storage Migration](#)

Prerequisites to Deploying in Azure

Overview

This article outlines the prerequisites for deployment of the Delphix Engine on Azure. The setup user should have experience launching and configuring instances in the Microsoft Azure environment. Review and complete the tasks in the next section before deployment.

Prerequisites

1. Review [Checklist of Information Required for Installation and Configuration](#).
2. Determine which virtual private cloud (VPC) is being used when launching the Virtualization instance. To maximize performance, deploy the Delphix Engine instance in the same VPC/subnet where the virtual databases (VDBs) will be created.

 Provisioning a VDB requires a compute instance running the same database engine as the source. Please note, however, that the target instance only needs storage to accommodate the OS, database platform binaries, etc., because Delphix delivers all of the data files.

3. Make sure that the necessary ports are open.

 Using the Delphix Engine for Azure will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform:

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network Access Requirements for SQL Server](#)

4. Update Security Group settings to accommodate the necessary connections.
 - a. Select the same **Network Security Group** that the current (or future) non-production Azure virtual machines utilize.
 - b. Modify the Network Security Group to allow access to all of the networking ports used by the Delphix Engine and the various source and target engines. See links above for information about specific port configurations.
5. Allocate storage.

 It is helpful to first create a list of the data sources intended for making dSources. A data source is typically a production database linked to the Virtualization Engine, enabling to create virtual, full, read-write copies of the source within minutes. The list should include the database name, platform (for example, Oracle or SQL Server), current size (in GB), the estimated number of virtual copies, and retention period (in days) of snapshots (backup copies).

- a. All data storage disks must be comprised of Azure Premium Storage devices. Delphix recommends using a minimum total of five disks to run the Delphix Engine. One disk is used for the boot device. The other four equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.
- b. Azure premium storage devices have different levels of guaranteed IOPs that vary from 120 IOPs to 20,000 IOPs. Please refer to [Microsoft's documentation](#) to ensure that the selected devices will meet the requirements for your deployment. UltraHD disk storage is supported for performance-sensitive customers.
- c. It is highly recommended to use Azure-managed disks.

- d. Use the guidelines outlined in [Deployment for Microsoft Azure](#).

Related Topics

- [Checklist of Information Required for Installation and Configuration](#)
- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network Access Requirements for SQL Server](#)
- [Deployment for Microsoft Azure](#)
- [Procedure for Deploying in Azure](#)

Procedure for Deploying in Azure

Overview

This article outlines the procedure for deploying the Delphix Engine in Microsoft Azure.

Deploying the Delphix Engine

1. Navigate to the Azure Marketplace at <https://azuremarketplace.microsoft.com>.
2. In the **Search** field, enter **Delphix**.
3. Select the **Delphix Virtualization for Azure (3TB)** for an hourly subscription or **Delphix DataOps Platform for Azure** for customers with existing licenses.
4. Review the information on the initial Marketplace page.
5. Click **GET IT NOW**.
6. Click **Create**.
7. Follow the creation wizard to deploy the Delphix Engine.
8. Provide the basic information for the Delphix Engine.
 - a. Select a name for the virtual machine.
 - b. For the OS disk type, select **Premium SSD**.
 - c. Enter a **username** and **password**.

 This information is never used but is part of the built-in wizard used by the marketplace. Any information provided in these fields is discarded.

- d. Select your **Subscription, Resource group, and Location**.
 - e. Click **OK**.
9. Select the **Size** of the virtual machine you want to deploy. For supported instance types and capabilities, See [Virtual Machine Requirements for Azure Platform](#).

 Each instance type will determine the limits for:

- Network bandwidth
- Maximum data disk
- Total IOPS

10. Click **Select**.
11. Configure the **Settings** for this virtual machine:
 - a. Use managed disks.
 - b. Select or create a **Virtual Network** and **Subnet**.
 - c. Optional: select a **Public IP address** (this will make an engine accessible over the internet)
 - d. Select or create a **Network Security Group**.
 - e. **Extensions** and **High Availability sets** are not currently supported. Please ensure that no extensions or availability sets have been configured.
 - f. Enable **Boot diagnostics**.
12. Click **OK**.
13. Review the configuration in the **Summary** screen and click **OK**.
14. Review the offer details in the **Buy** screen and click **Purchase** to deploy the Delphix Engine.

Configuring the Delphix Engine Virtual Machine

1. From the **Azure Portal** sidebar, select **Virtual Machines**. The running virtual machines will be displayed.
2. Click the virtual machine's **name** to open its detailed information.

3. Click **Disks** to open the detailed panel.
4. Click **Add data disk**.
5. Create a new device:
 - a. In the **Create disk** drop-down menu, select **Create managed disk**. This will open the managed disk wizard.
 - b. In the wizard, ensure that the **Account type** indicates **Premium (SSD)** and that the **Source Type** shows **None (empty disk)**.
 - c. Select the size of the device to create. Note that the device size will determine the IOPS limit for that device.
 - d. Click **Create**.
6. Once the device is created, change the **Host Caching** parameter for that device to **None**.
7. Repeat the above steps to create additional disks.
8. Once adding all the disks is complete, click **Save**.
9. Optional: modify the virtual machine's **Network Security Group**.
 - a. From the side panel of the virtual machine's detail screen, click **Network interfaces**.
 - b. Click the name of the network interfaces to modify.
 - c. From the **Network interface** details window, click the **Network security group** specified in the **Overview** pane. In the example below, the user would click **delphix-engine-nsg**:
 - d. From the **Network security group** details pane, modify the **Inbound security rules** and/or the **Outbound security rules** by selecting from the side panel and clicking **Add**.
 - e. Add the appropriate rule and click **OK**. The screenshot below shows adding the inbound rule for the Delphix Session Protocol.
 - f. Repeat the above steps for all the rules that should be added.

Configuring the Delphix Engine

1. Connect to the running Delphix instance with a web browser. Use the Public IP address or DNS name noted in the **Network Interface** Panel. Upon successful connection, the browser will automatically redirect to the **Delphix Setup Page**.
2. Refer to the standard [product deployment instructions](#) to complete the Delphix deployment.

Next Steps

Congratulations! The Delphix Virtualization Engine should be successfully deployed in Azure.

Use Delphix documentation to learn how to:

- configure a database source
- configure the target environments
- create virtual databases (VDBs)

Related Topics

- [Deployment for Microsoft Azure](#)
- [Prerequisites to Deploying in Azure](#)

Enabling Azure Accelerated Networking

Overview

This article describes how to enable Azure Accelerated Networking. If a VM was created without Accelerated Networking, enabling this feature on an existing VM is possible. The VM must meet the following prerequisites:

- VM must be a supported size for Accelerated Networking
- VM must be a supported Azure Gallery image (and kernel version for Linux)
- All individual VMs or VMs in an availability set must be stopped/deallocated before enabling Accelerated Networking on any NIC

For more information, please see the [Azure documentation](#).

Individual VMs and VMs in an Availability Set

Stop/deallocate the VM or, if an Availability Set, all the VMs in the Set:

- i**
- If the VM was created individually (without an availability set), only the individual VM needs to stop/deallocate to enable Accelerated Networking.
 - If the VM was created with an availability set, all VMs contained in the availability set will need to stop/deallocate before enabling Accelerated Networking on any Network Interface Card (NIC).

Azure CLI

```
az vm deallocate \
  --resource-group myResourceGroup \
  --name myVM
```

Once stopped, enable Accelerated Networking on the NIC of the VM:

Azure CLI

```
az network nic update \
  --name myNic \
  --resource-group myResourceGroup \
  --accelerated-networking true
```

Restart the VM (or all VMs if using an availability set) and confirm that Accelerated Networking is enabled:

Azure CLI

```
az vm start --resource-group myResourceGroup \
```

```
--name myVM
```

After the restart, the Mellanox VF (Virtual Function) device will be exposed to the VM.

Deployment for Google Cloud Platform

Overview

This article covers the virtual machine requirements, including memory and data storage, for the deployment of the Delphix Engine on Google Cloud Platform (GCP). Once the requirements listed on this page are reviewed, refer to the next articles on GCP deployment:

- [Prerequisites to Deploying in GCP](#)
- [Procedure for Deploying in GCP](#)

Machine Types

The following is a list of instance types that are supported to deploy Delphix on GCP. Delphix periodically certifies new instance types, which will be added to the list here.

Recommendations	Notes
<ul style="list-style-type: none"> • N2 standard • N2 high-mem 	<ul style="list-style-type: none"> • The Delphix Engine most closely resembles a storage appliance and performs best when provisioned using a storage-optimized instance type. • Larger instance types provide more CPU, which can prevent resource shortfalls under high I/O throughput conditions. • Larger instances also provide more memory, which the Delphix Engine uses to cache database blocks. More memory will provide better read performance. • Recommended minimum size of 8 vCPU and 64G memory.

Network Configuration

Requirements	Notes
Virtual Private Cloud	<p>The Delphix Engine and all of the source and target environments must be deployed in a VPC network to ensure that private IP addresses are static and do not change when restarting instances.</p> <p>When adding environments to the Delphix Engine, use the host's VPC (static private) IP addresses.</p>
Static Public IP	<p>The GCP Delphix instance must be launched with a static IP address; however, the default behavior for VPC instances is to launch with a dynamic public IP address – which can change whenever the instance restarts. If using a public IP address for the Delphix Engine, static IP addresses can be assigned using the Google Cloud documentation.</p>

Security Group Configuration	<p>The default security group will only open port 22 for SSH access. The security group must be modified to allow access to all of the networking ports used by the Delphix Engine and the various source and target engines.</p> <p>See Network Performance Configuration Options for information about network performance tuning.</p> <p>See General Network and Connectivity Requirements for information about specific port configurations.</p>
------------------------------	---

Storage Configuration

Requirements	Notes
General Storage Configuration	<p>Allocate initial storage equal to the size of the physical source database storage. For high redo rates and/or high DB change rates, allocate an additional 10-20 %.</p> <p>Add storage when storage capacity approaches 30% free.</p> <p>Delphix recommends using a minimum of four equally sized storage disks to run the Delphix Engine. This allows the Delphix File System (DxFS) to assure its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.</p>

Related Topics

- [Checklist of Information Required for Installation and Configuration](#)
- [Prerequisites to Deploying in GCP](#)
- [Procedure for Deploying in GCP](#)

Prerequisites to Deploying in GCP

Overview

This article outlines the prerequisites for deployment of the Delphix Engine on GCP. The setup user should have experience launching and configuring instances in the Google Cloud Platform environment. Review and complete the tasks in the next section before deployment.

Prerequisites

1. Review the [Checklist of Information Required for Installation and Configuration](#).
2. Make sure that the Google account being used to deploy the Delphix Engine has an appropriate level of enablement for the **Delphix Virtualization for GCP** subscription.
3. Determine which virtual private cloud (VPC) is being used when launching the virtualization instance. To maximize performance, deploy the Delphix Engine instance in the same VPC/subnet in which the virtual databases (VDBs) will be created.

 Provisioning a VDB requires a compute instance running the same database engine as the source. Please note, however, that the target instance only needs storage to accommodate the OS, database platform binaries, etc., because Delphix delivers all of the data files.

4. Make sure that the necessary ports are open.

 Using the Delphix Engine for GCP will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform.

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network and Connectivity Requirements for SQL Server Environments](#)
- [Network and Connectivity Requirements for SAP ASE Base Environments](#)
- [Network and Connectivity Requirements for Db2 Environments](#)

5. Update Security Group settings to accommodate the necessary connections.
 - a. Select the same Security Group that the current (or future) non-production EC2 compute nodes utilize.
 - b. Modify the Security Group to allow access to all of the networking ports used by the Delphix Engine and the various source and target platforms. See links above for information about specific port configurations.
6. Allocate storage.
 - a. To properly size the initial storage capacity and determine the number and size of Provisioned IOPs Volumes required, download and utilize the [Delphix-Dynamic-Data-Platform-Storage-Calculator](#).

 It is helpful to first create a list of the data sources intended for making dSources. A data source is typically a production database linked to the Virtualization Engine, enabling to create virtual, full, read-write copies of the source within minutes. The list should include the database name, platform (for example, Oracle or SQL Server), current size (in GB), the estimated number of virtual copies, and retention period (in days) of snapshots (backup copies).

- b. Delphix recommends using a minimum total of four disks to run the Delphix Engine. One disk is used for the boot device. The other four equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.

Additional GCP Configuration Notes

- Delphix supports both Zonal and Regional SSD persistent disks.

Related Topics

- [Deployment for GCP](#)
- [Procedure for Deploying in GCP](#)

Procedure for Deploying in GCP

Overview

This article outlines the procedure for deploying the Delphix Engine using a tar.gz file in Google Cloud Console or in Google Cloud Marketplace.

Deployment in Google Cloud Console

1. Download the latest **Delphix Platform for GCP** tar.gz file from <https://download.delphix.com/>.
2. Create a **GCP storage bucket** and upload the tar.gz file from the first step.
3. Once the upload is complete, navigate to **Compute Engine > Images** and select **Create Image**.
4. In the **Create Image** dialog, provide the following information.
 - a. Image Name
 - b. Source: Cloud Storage File
 - c. Select the bucket and newly uploaded Delphix File.
 - d. Location: Multi-Regional or Regional. (To deploy Delphix in multiple regions - select **Multi-Region**)
 - e. Family: Optional
 - f. Description: Optional
 - g. Labels: Optional
 - h. Encryption: Select based on the organization's policy
 - i. Click **Create**. (Creating the image could take some time)
5. Once the image creation is complete, select the newly created image and create an instance from it.
6. Configure the instance creation screen with the following items.
 - a. Identify and select the region and zone to run Virtual Databases.
 - b. Choose a supported machine type, commensurate with the expected workload.
 - c. Change/modify the **Boot Disk** to **SSD Persistent Disk**. (Defaults to **Standard Persistent Disk**)
 - d. Add **Data Disks** (minimum of four recommended) with a total size at least equal to the total source DB sizes. Be sure to utilize SSD Persistent Disks.
 - e. Update Networking, commensurate with the Network and Subnet(s) on which the target non-production instances reside. Note, a “target” is an instance running the DB platform identical to the source DBs.
 - f. Once all of the necessary instance information has been configured, select **Create**.
 - g. Wait for the instance to be created/available and connect to the newly deployed Engine using the assigned IP/hostname via the support web browser.

Deployment in Google Cloud Marketplace

1. Login to the **Google Cloud Marketplace**.
2. In the **Search** field, enter **Delphix**.
3. Select the **Delphix Data Virtualization for GCP (3TB)**.
4. Review the information on the initial Marketplace page.
5. Click **Launch**.
6. Review the deployment configuration and software subscription information.
7. Accept Google Cloud’s and Delphix Corp’s Terms of Service.
8. Click **Deploy**, which will start the deployment of the instance.
9. After the deployment, add equally-sized data disks to the instance.

Configuring the Delphix Engine

1. Connect to the running Delphix instance with a web browser. Use the IP address or DNS name noted in the Instance Description. Upon successful connection, the browser will automatically redirect to the **Delphix Setup Page**.
2. Refer to the standard [product deployment instructions](#) to complete the Delphix deployment.

Next Steps

Congratulations! The Delphix Virtualization Engine should be successfully deployed in Google Cloud Platform.

Use Delphix documentation to learn how to:

- [configure a database source](#)
- [configure the target environments](#)
- [configure virtual databases \(VDBs\)](#)

Related Topics

- [Deployment for Google Cloud Platform](#)
- [Prerequisites to Deploying in GCP](#)

Deployment for OCI

Overview

This article outlines the virtual machine requirements, including memory and data storage, for deploying the Delphix Engine on Oracle Cloud Infrastructure (OCI). Once the requirements listed on this page are reviewed, refer to the next articles on OCI deployment:

- [Prerequisites to Deploying in OCI](#)
- [Procedure for Deploying in OCI](#)

Compute Image Types

Delphix distributes product images for OCI using the QCOW2 image type. Compute Images must be imported into OCI using the Paravirtualized launch mode; currently, images using the Emulated launch mode are not supported.

Supported Shapes

The following is a list of shapes that are supported to deploy Delphix on OCI.

Requirements	Notes
Large Memory Instances (preferred) <ul style="list-style-type: none"> • VM.Standard2 • VM.Standard.E4.Flex 	<ul style="list-style-type: none"> • The Delphix Engine most closely resembles a storage appliance and performs best when provisioned using a storage-optimized shape. • Larger shapes provide more CPU, which can prevent resource shortfalls under high I/O throughput conditions. • Larger shapes also provide more memory, which the Delphix Engine uses to cache database blocks. More memory will provide better read performance. • Recommended minimum size of 8 vCPU and 64G memory.

Network Configuration

Requirements	Notes
Virtual Cloud Network (VCN)	<ul style="list-style-type: none"> • The Delphix Engine and all of the source and target environments must be deployed in a VCN to ensure that private IP addresses are static and do not change when restarting instances. • By default, OCI subnets are considered public. When defining a subnet, it is recommended to set it as private. Unless required by the environment, the VCN should not include a Public Subnet. • When adding environments to the Delphix Engine, use the host's VCN (static private) IP addresses.

Requirements	Notes
Static Private IP	<ul style="list-style-type: none"> The Delphix instance should be launched with a static private IP address. For security reasons, it is encouraged to avoid configuring the engine with a Public IP address; however, it could be passable to use a dynamic Public IP address in addition to a static Private IP address if the environment requires such access.
Security Rules Configuration	<ul style="list-style-type: none"> OCI allows two firewall features: Network Security Groups (NSGs) and Security Lists. Oracle recommends the use of NSGs over Security Lists because NSGs let the VCN subnet architecture be separate from the application security requirements. More information can be found in this Oracle documentation. A VCN will use a Security List to define default rules. By default, the security list will only open port 22 for SSH access. The Security List must be modified, or new NSGs created, to allow access to all of the networking ports used by the Delphix Engine and the various source and target engines. This dual implementation of firewall or security rules may be different from other clouds, please see OCI documentation for best practices. See Network Performance Configuration Options for information about network performance tuning. See General Network and Connectivity Requirements for information about specific port configurations.

General Storage Configuration

Requirements	Notes
General Storage Configuration	<ul style="list-style-type: none"> Allocate initial storage equal to the size of the physical source database storage. For high redo rates and/or high DB change rates, allocate an additional 10-20 %. Add storage when storage capacity approaches 30% free. Delphix recommends using a minimum of four equally sized storage disks to run the Delphix Engine. This allows the Delphix File System (DxFS) to assure its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage. Must use Block Volume for data storage. Block Volumes must be attached using Paravirtualized mode.

Requirements	Notes
OCI Storage Configuration	<ul style="list-style-type: none"> • Currently supported Shapes only support Block Volumes; File Storage is not supported. • Paravirtualized block devices are required; currently, iSCSI devices are not supported. • Elastic Performance Configuration Options (aka Volume Performance Policy): use Higher Performance.

Additional OCI Configuration Notes

- When running low on storage space, Delphix recommends adding additional equivalently sized block storage volumes, or devices, instead of resizing existing volumes.
- If existing storage volumes must be expanded, this must be done using the “online” resizing strategy specified in OCI documentation; “offline” storage resizing is not supported and may lead to unexpected downtime. If an existing storage volume is expanded, use the Setup or sysadmin interface to expand each storage “device” or volume. The additional storage, as a result of a resize, will not be available for use until the storage devices are explicitly instructed to make use of the additional space.
- If expanding storage volumes, it is recommended that all volumes are expanded to the same size. When storage volumes or devices are the same size, the Delphix product is able to balance I/O distribution among the disks for optimal performance.
- Hot removal of storage volumes is not supported.

Related Topics

- [Network Connectivity Requirements](#)
- [Prerequisites to Deploying in OCI](#)
- [Procedure for Deploying in OCI](#)
- [Oracle Matrix](#)

Prerequisites to Deploying OCI

Overview

This article outlines the prerequisites for deployment of the Delphix Engine on OCI. The setup user should have experience launching and configuring instances in the Oracle Cloud Infrastructure environment. Review and complete the tasks in the next section before deployment.

Prerequisites

1. Contact a Delphix representative to request this capability. Delphix will assist in assuring that all Oracle Cloud requirements are met to successfully run a Delphix Engine with the most appropriate configuration for the use case.
2. Review the [Checklist of Information Required for Installation and Configuration](#).
3. Review Oracle's [Creating an Instance documentation](#).
4. Make sure that the Oracle account being used to deploy the Delphix Engine has an appropriate level of permissions to upload files to a Bucket in Object Storage, and use that object to produce a Custom Image.
5. Determine which virtual cloud network (VCN) is being used when launching the virtualization instance. To maximize performance, deploy the Delphix Engine instance in the same VCN/subnet in which the virtual databases (VDBs) will be created.

 Provisioning a VDB requires a compute instance running the same database engine as the source. Please note, however, that the target instance only needs storage to accommodate the OS, database platform binaries, etc., because Delphix delivers all of the data files.

6. Make sure that the necessary ports are open.

 Using the Delphix Engine for OCI will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform.

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network and Connectivity Requirements for SQL Server Environments](#)
- [Network and Connectivity Requirements for SAP ASE Environments](#)
- [Network and Connectivity Requirements for Db2 Environments](#)

7. Update Security Group (or Security List) settings to accommodate the necessary connections.

1. Select the same Security Group that your current (or future) non-production compute nodes utilize.
 - b. Modify the Security Group to allow access to all of the networking ports used by the Delphix Engine and the various source and target platforms. See [Oracle Network Security Groups](#).
- Allocate storage.
 - a. To properly size the initial storage capacity and determine the number and size of Provisioned IOPs Volumes required, download and utilize the [Delphix-Dynamic-Data-Platform-Storage-Calculator](#).

 It is helpful to first create a list of the data sources intended for making dSources. A data source is typically a production database linked to the Virtualization Engine, enabling to create virtual, full, read-write copies of the source within minutes. The list should include

the database name, platform (for example, Oracle or SQL Server), current size (in GB), the estimated number of virtual copies, and retention period (in days) of snapshots (backup copies).

- b. Delphix only supports using Block Volume devices for data storage. These block devices must be attached to the Delphix Engine using the Paravirtualized attach mode (iSCSI is not currently supported). Block Volumes must be provisioned in the same Availability Domain as the Delphix Engine, otherwise, they will not be discoverable. Block Volumes should be configured to use the “[Higher Performance](#)” performance option. All Block Volumes attached to a single Delphix Engine should be of the same size. A minimum of four (4) Block Volumes should be attached to a Delphix Engine.
- During the Manual Deployment option, use the guidelines outlined in [Virtual Machine Requirements for OCI](#).

Related Topics

- [Deployment for OCI](#)
- [Checklist of Information Required for Installation and Configuration](#)
- [Procedure for Deploying in OCI](#)

Procedure for Deploying in OCI

Overview

This article outlines the procedure for deploying the Delphix Engine in Oracle Cloud Infrastructure.

Download and Verify the Delphix Engine Image

1. Contact an account manager to request access to the OCI variant of the Delphix product.
2. Follow the link given by the Delphix solutions architect. Download the **Delphix_6.x.x.x_....Standard_OCI.qcow2** file and the **SHA256SUMS** file.
3. Once both files have finished downloading and are in the same directory, run the following command to verify the download:

```
$ grep -i OCI.qcow2 ./SHA256SUMS | sed -E 's,Appliance_Images/
(Controlled_Availability/)?, ,g' | sha256sum --check
```

Upload the Delphix Engine Image as an Object

1. Authenticate with OCI and navigate to the [Infrastructure Console](#).
2. Use the navigation menu to reach the [Object Storage Buckets, Core Infrastructure, page](#) (Menu > Object Storage > Object Storage).
3. Set the **List Scope Compartment**. This will depend on the organization's strategy for managing OCI resources.
4. [Create a storage bucket](#) or select an existing bucket.
5. Click the blue **Upload** button.
6. In the **Upload Objects** modal window, specify an optional prefix and choose the OCI specific QCOW2 file that was previously downloaded.
7. Click the blue **Upload** button.

Creating a Custom Compute Image from an Object

1. Authenticate with OCI and navigate to the [Infrastructure Console](#).
2. Use the navigation menu to reach the [Compute Custom Images, Core Infrastructure, page](#) (Menu > Compute > Custom Images).
3. Set the **List Scope Compartment**.
4. Click the blue **Image Import** button.
5. In the **Import Image** modal window, select a suitable compartment in the **Create In Compartment** field that conforms to the organization's strategy on managing OCI resources.
6. In the **Name** field enter a unique name to identify the Custom Compute Image (using the same name as the image object from the previous step is recommended). Upload the Delphix Engine Image as an Object.
7. For **Operating System** select **Linux**.
8. Next, identify an object by specifying its Compartment, Bucket, and Object Name. Or, specify an Object Storage URL.

 The Object Details will identify this value as **URL Path (URI)**.

9. For **Image Type** select **QCOW2**.
10. For **Launch Mode** select **Paravirtualized Mode**.
11. For organizations that have a tagging policy for cloud-based resources, expand the **Tagging Options** section, and define tags.

- Click the blue **Import Image** button.

Launching the Delphix Engine

- Authenticate with OCI and navigate to the [Infrastructure Console](#).
- Use the navigation menu to reach the [Compute Instances, Core Infrastructure, page](#) (Menu > Compute > Instances).
- Set the **List Scope Compartment**.
- Click the blue **Create Instance** button.
- In the **Create Compute Instance** window pane, specify a unique name for the VM.
- For the **Create In Compartment** field, select a suitable compartment that conforms to the organization's strategy on managing OCI resources.
- In the **Image or operating system** section, click the **Change Image** button. Switch to the **Custom Images** tab. Find the Delphix image that corresponds to the instance being deployed. Click the blue **Select Image** button.

 If the Delphix Custom Image is not visible, look for the **Change Compartment** option near the top of the current window pane.

- Each Availability Domain has its own quota, it is ok to use AD-1, AD-2, or AD-3 - but, be sure to make note of which Availability Domain is being used.

 Compute Instances and attached Storage will need to be in the same Availability Domain.

- In the **Shape** section, click the **Change Shape** button. For **Instance type**, specify **Virtual Machine**. For **Shape series**, use **Intel Skylake**. Select an OCI Shape that is supported by Delphix.
- Continue on to the **Configure networking** section.
 - If the network is not specified correctly, it is likely that firewall issues will occur. In this case, please consult the organization's IT or DevOps team. If the organization is using Network Security Groups (NSGs), mark the **Use Network Security Groups to Control Traffic** checkbox after consulting the appropriate teams. Last, select the **Do Not Assign a Public IP Address** radio button.
- The **Boot Volume** section can be skipped.
- In the **Add SSH Keys**, select the **No SSH Keys** radio option. The Delphix product is a closed appliance and manages users independently.
- In general, all of the Advanced Options can be skipped. For organizations that have a tagging policy for cloud-based resources, expand into the Advanced Management section, and look for the **Tagging** sub-section to define tags.
- Click the blue **Create** button - wait about 2-5 minutes for the Delphix Engine instance to boot.

Create Block Storage Volumes

- Authenticate with OCI and navigate to the [Infrastructure Console](#).
- Use the navigation menu to reach the [Block Volumes, Core Infrastructure, page](#) (Menu > Block Storage > Block Volumes).
- Set the **List Scope Compartment**.
- Click the blue **Create Block Volume** button.
- In the **Create Block Volume** modal window, specify a unique name for this Block Volume. It can be helpful if this name is descriptive or identifies the VM it is intended to be attached to and ends in a sequence number.
- For the **Availability Domain**, this value **MUST** be the same Availability Domain used for the Delphix Engine instance, otherwise, this volume will not be available for use.
- In the **Volume Size and Performance** section, select the **Custom** option. Set the size of the volume to be sufficiently large (with room for growth) to support the databases that will be virtualized, or masked, by this Delphix Engine. Set the **Default Volume Performance** to the **Higher Performance** setting.

8. A **Backup Policy** is not required and can be left blank or **No Backup Policy Selected**. However, depending on the organization's needs, consider selecting a Backup Policy.
9. For **Encryption**, the default option of **Encrypt Using Oracle-Managed Keys** is permissible. Optionally, for independent encryption keys, use the **Encrypt Using Customer-Managed Keys** option.
10. For organizations that have a tagging policy for cloud-based resources, expand the **Tagging Options** section and define tags.
11. Uncheck the checkbox that says **View Detail Page After This Block Volume is Created**. This will prevent navigating away from the Block Volumes page because multiple Block Volumes will need to be created at the same time.
12. Click the blue **Create Block Volume** button.
13. A Delphix Engine requires a minimum of four equally sized Block Volumes for data storage. Repeat Steps 4-12 as many times as necessary.

Attach Block Storage Volumes

1. Authenticate with OCI and navigate to the [Infrastructure Console](#).
2. Use the navigation menu to reach the [Block Volumes, Core Infrastructure, page](#) (Menu > Block Storage > Block Volumes).
3. Set the **List Scope Compartment**.
4. From the list of pre-existing Block Volumes, identify the resources being attached to the Delphix Engine and wait until the volume state becomes **Available**.
5. Select one of the **Block Volumes** to enter the **Block Volume Details** page.
6. On the left-hand side, locate the **Resources** menu and select **Attached Instances**.
7. If the Block Volume has not been previously attached to another VM, selecting the blue **Attach to Instance** button will be available.
8. In the **Attach to Instance** modal window, specify the **Attachment Type** as **Paravirtualized**. Currently, iSCSI is not supported.
9. For **Access Type** use the **READ/WRITE** option.
10. Next, identify a Delphix Engine by selecting an instance or by specifying an instance OCID. If the Delphix Engine instance is not shown in the **Select an Instance** drop-down menu, the **Change Compartment** option could be needed. Block Volumes can only be attached to VM instances that were created in the same Availability Domain - if these values do not match, either re-provision Block Volumes or the Delphix Engine in the correct Availability Domain.
11. Click the blue **Attach** button.
12. Repeat Steps 4-11 until all associated Block Volume resources have been attached to the Delphix Engine instance.

Configuring the Delphix Engine

1. Connect to the running Delphix Engine instance with a web browser. Use the IP address or FQDN noted in the Instance Description.
Upon successful connection, the browser will display a login prompt to enter the Delphix Setup Page.
2. Refer to the standard [product deployment instructions](#) to complete the Delphix deployment.

Next Steps

Congratulations! The Delphix Virtualization Engine should be successfully deployed in OCI.

Use Delphix documentation to learn how to:

- configure a database source
- configure target environments
- create virtual databases (VDBs)

Related Topics

- [Deployment for OCI](#)
- [Prerequisites to Deploying in OCI](#)

Deployment for IBM Cloud

Overview

This article covers the virtual machine requirements, including memory and data storage, for the deployment of the Delphix Engine on IBM Cloud. Once the requirements listed on this page are reviewed, refer to the next articles on IBM Cloud deployment:

- [Prerequisites for Deploying in IBM Cloud](#)
- [Procedure for Deploying in IBM Cloud](#)

Supported Profiles

The following is a list of profiles that are supported to deploy Delphix on IBM Cloud.

Requirements	Notes
<ul style="list-style-type: none"> • mx2 	<ul style="list-style-type: none"> • The Delphix Engine most closely resembles a storage appliance and performs best when provisioned using a storage-optimized profile. • Larger profiles provide more CPU, which can prevent resource shortfalls under high I/O throughput conditions. • Larger profiles also provide more memory, which the Delphix Engine uses to cache database blocks. More memory will provide better read performance. • Recommended minimum size of 8 vCPU and 64G memory.

Network Configuration

Requirements	Notes
Virtual Server Instances	<ul style="list-style-type: none"> • Deploy the Delphix Engine and all of the source and target environments in the same VPC network. • When adding environments to the Delphix Engine, use the host's VPC IP addresses.
Security Configuration	<ul style="list-style-type: none"> • The default security group will only open port 22 for SSH access. The security group must be modified to allow access to all of the networking ports used by the Delphix Engine and the various source and target engines. • See Network Performance Configuration Options for information about network performance tuning. • See General Network and Connectivity Requirements for information about specific port configurations. • Reference: IBM Cloud Security and Compliance documentation

Port Requirements

Using the Delphix Engine for IBM Cloud will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform.

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network and Connectivity Requirements for SQL Server Environments](#)
- [Network and Connectivity Requirements for SAP ASE Environments](#)
- [Network and Connectivity Requirements for Db2 Environments](#)

General Storage Configuration

Requirements	Notes
General Storage Configuration	<ul style="list-style-type: none"> • Allocate initial storage equal to the size of the physical source database storage. For high redo rates and/or high DB change rates, allocate an additional 10-20 %. • Add storage when storage capacity approaches 30% free. • Delphix recommends using a minimum of four equally sized storage disks to run the Delphix Engine. This allows the Delphix File System (DxFS) to assure its file systems are always consistent on disk without additional serialization. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage. • Reference: IBM Block Storage Documentation.

Additional IBM Configuration Notes

- Expandable volume is a beta feature that is available for evaluation and testing purposes. This feature is available in the US South, US East, London, and France regions. Contact an IBM Sales representative if interested in getting access at <https://cloud.ibm.com/docs/vpc?topic=vpc-expanding-block-storage-volumes>.
- After performing an “online” resize/expansion of a storage volume using IBM Cloud tools, then use the Delphix sysadmin interface to “Expand” the storage device; otherwise, the newly allocated storage space, from the resize/expansion, will not be used.
- Resize/expansion of a storage volume using IBM Cloud is not supported while the Delphix engine is in a stopped state.
- Removing a storage volume should be done while the machine is running. First use the Delphix sysadmin CLI interface to “Unconfigure” the storage device, then remove it from IBM Cloud.

Related Topics

- [Network Connectivity Requirements](#)
- [Procedure for Deploying in the IBM Cloud](#)
- [Support Matrix](#)

Prerequisites for Deploying in IBM Cloud

Overview

This article outlines the prerequisites for deployment of the Delphix Engine on IBM Cloud. The setup user should have experience launching and configuring instances in the IBM Cloud environment. Review and complete the tasks in the next section before deployment.

Prerequisites

1. A Delphix software license is required. New users should [contact Delphix](#) to get started.
2. Review the [Checklist of Information Required for Installation and Configuration](#).
3. Review IBM's [cloud documentation](#) for IBM Cloud-specific information.
4. Determine which virtual private cloud (VPC) is being used when launching the virtualization instance. To maximize performance, deploy the Delphix Engine instance in the same VCN/subnet in which the virtual databases (VDBs) will be created.

 Provisioning a VDB requires a compute instance running the same database engine as the source. Please note, however, that the target instance only needs storage to accommodate the OS, database platform binaries, etc., because Delphix delivers all of the data files.

5. Make sure that the necessary ports are open.

 Using the Delphix Engine for GCP will require connections to source and target database servers. Such connections require various ports to be open, enabling communications. For a detailed list of the network and port requirements, click the link that corresponds with the relevant database platform.

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Network Access Requirements for SQL Server](#)
- [Network and Connectivity Requirements for SAP ASE Environments](#)
- [Network and Connectivity Requirements for Db2 Environments](#)

6. Update Security Group settings to accommodate the necessary connections.
7. Select the same Network Security Group that your current (or future) non-production Azure virtual machines utilize.
8. Modify the Security Group to allow access to all of the networking ports used by the Delphix Engine and the various source and target platforms. See links above for information about specific port configurations.
9. Allocate storage.

 It is helpful to first create a list of the data sources from which you intend to make dSources. A data source is typically a production database that you link to the Delphix Engine, enabling you to create virtual, full, read-write copies of the source within minutes. The list should include the database's name, platform (for example, Oracle or SQL Server), current size (in GB), the estimated number of virtual copies, and retention period (in days) of snapshots (backup copies).

- a. All data storage disks must be comprised of Premium Storage devices.
- b. Delphix recommends using a minimum total of four disks to run the Delphix Engine. One disk is used for the boot device. The other four equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage.

- c. Premium storage devices have different levels of guaranteed IOPs that vary from 120 IOPs to 7500 IOPs. Please refer to IBM Cloud documentation to ensure that the selected devices will meet the requirements for your deployment.
- d. Use the guidelines outlined in Procedure for Deploying in IBM Cloud.

Related Topics

- [Deployment for IBM Cloud](#)
- [Procedure for Deploying in the IBM Cloud](#)

Procedure for Deploying in the IBM Cloud

Overview

This article outlines the procedure for deploying the Delphix Engine in IBM Cloud. There are two methods for deploying a Delphix Engine in the IBM Cloud using the [Software Catalog](#) or [Manually Uploading the Delphix Image](#).

Deploying from the IBM Software Catalog

1. Navigate to the [IBM Software Catalog](#) and search for Delphix.
2. Select the Delphix Data Virtualization Tile or Masking Tile for the Masking product.
3. Scroll down to the Deployment Values section and input specifics for the environment.

Required Parameters	Description
hostname	The name of the VSI used to deploy Delphix.
profile	Compute profile to be used for deploying Delphix (see recommended profiles).
ssh_key	Public SSH key is to be used when provisioning the VSI.
subnet_id	The id of the subnet where the VSI will be provisioned.
volumecount	Number of block storage volumes
volumeprofile	Block storage profile to use (recommended is >= 10 IOPS/GB)
volumesize	Block storage volume size.
vpcname	Name of the VPC where the VSI is provisioned.
zone	VPC zone to provision the environment.

Downloading the Delphix Image

 Contact an account manager to request access to the IBM variant of the Delphix product.

1. Follow the link given by the Delphix solutions architect. Download the **Delphix_Verson...._Standard_IBM.qcow2** file and the **SHA256SUMS** file.
2. Once both files have finished downloading and are in the same directory, run the following command to verify the download:

```
$ grep -i IBM.qcow2 ./SHA256SUMS | sed -E 's,Appliance_Images/
(Controlled_Availability)?, ,g' | sha256sum --check
```

Uploading the Delphix Engine Image as an Object

1. Authenticate with the IBM Cloud and navigate to the [Dashboard](#).
2. Use the navigation menu to reach the [Resource List page](#). The Resource List page can be navigated from the Dashboard by clicking on **Storage** within the Resource Summary pane.
3. Expand Storage from the menu and select the appropriate resource group. Depending on the organization's strategy for managing IBM resources, [creating a resource group](#) may be required.
4. [Create a storage bucket](#) or select an existing bucket.
5. Click the blue **Upload** button and select **Files**.
6. When the pop-up menu appears, select the **Transfer Type**. Aspera High-Speed Transfer is required for large files, the plugin will need to be installed. It will automatically navigate through the steps to install the plugin.
7. In the **Upload Files** (objects) window, click on the **Select Files** (objects) button and choose the **IBM-specific QCOW2** file that was previously downloaded.
8. Click the **Upload** button.

Creating a Custom Image

1. Authenticate with IBM Cloud and navigate to the [Dashboard](#).
2. Use the navigation menu to reach the [Custom images page](#) for VPC within the VPC infrastructure (IBM Cloud pull-down menu > VPC Infrastructure > Custom images).
3. Click the blue **Create** button.
4. In the **Import Custom Image** page, specify a unique name for the image.
5. From the **Resource Group** drop-down, select the organization's resource group.
6. Optional: In the **Tags** section, provide appropriate [tags](#) to organize resources.
7. Select the appropriate **Region**.
8. Select the **Cloud Object Storage** bucket containing the uploaded image by navigating to **Instances > Location > Bucket** from the drop-down menus. The downloaded **QCOW2** image should appear in the pane below the three drop-down menus.
9. Within the **Operating System** section, click on the **Ubuntu Linux** tile and select **ubuntu-18-04-amd64** from the drop-down menu.
10. Once all the parameters are entered, in the right pane click on the blue button to import custom image.

Launching the Delphix Engine

1. Authenticate with IBM Cloud and navigate to the [Dashboard](#).
2. Use the navigation menu to reach the [Virtual Server Instances page](#) within the VPC Infrastructure (IBM Cloud pull-down menu > VPC Infrastructure > Virtual Server Instances). Note: To maximize performance, deploy the Delphix Engine instance in the same VPC/subnet in which you will create your virtual databases (VDBs).
3. Click the blue **Create** button.
4. In the **New Virtual Server** for VPC page, specify a unique name for the VM.
5. From the **Virtual Private Cloud** drop-down, select the organization's VPC.
6. From the **Resource Group** drop-down, select the organization's resource group.
7. Optional: In the **Tags** section, provide appropriate [tags](#) to organize resources.
8. Select the **Location** of the IBM Cloud resources.
9. In the **Operating System** section, click on the **Select Custom Image** link within the Custom Image block.
10. In the pop menu, select the IBM-specific image previously uploaded.
11. Within the **Profile** section, click on **View** all profiles. Select one of the supported instance types and click **Save**.
12. The **User Data** section can be skipped.

13. The Boot Volume section can be skipped since it would already have the default values.
14. Creating block storage volumes can be done at a later time and will be discussed in the next section.
15. Continue on to the **Network Interfaces** section. If a subnet is already configured in the zone and VPC, then this section will already have a default network interface. Otherwise, create a subnet with the appropriate security groups. If the network is not specified correctly, it is likely that firewall issues will occur. In this case, please consult the organization's IT or DevOps team. If the organization is using Network Security Groups (NSGs), mark the **Use Network Security Groups to Control Traffic** checkbox after consulting the appropriate teams. Last, select the **Do Not Assign a Public IP Address** radio button.
16. Click the **Create virtual server instance** button on the right panel. This will take a couple of minutes.

Creating Block Storage Volumes

1. Authenticate with IBM Cloud and navigate to the [Dashboard](#).
2. Use the navigation menu to reach the [Block Storage Volumes](#) within VPC Infrastructure (IBM Cloud pull-down menu > VPC Infrastructure > Block Storage Volumes).
3. Click the blue **Create** button.
4. In the **Block Storage Volume for VPC** modal window, specify a unique name for this Block Volume. It can be helpful if this name is descriptive or identifies the VM it is intended to be attached to and ends in a sequence number.
5. From the **Resource Group** drop-down, select the organization's resource group.
6. Optional: In the **Tags** section, provide appropriate [tags](#) to organize resources.
7. Select the **Location** of the IBM Cloud resources.
8. Enter the required **IOPS**. The recommended supported IOPS is 10/GB.
9. Enter the storage size in GB. Set the size of the volume to be sufficiently large, with room for growth, to support the databases that will be virtualized, or masked, by this Delphix Engine.
10. Keep the **Encryption** settings at default, e.g. Provider Managed.
11. Click the blue **Create Volume** button.
12. Delphix recommends using a minimum total of four disks to run the Delphix Engine. One disk is used for the boot device. The other four equally sized disks will be used for data storage. This also enables the Delphix Engine to achieve higher I/O rates by queueing more I/O operations to its storage. Repeat Steps 3-11 as many times as necessary.

Attaching Block Storage Volumes

1. Authenticate with IBM Cloud and navigate to the [Dashboard](#).
2. Use the navigation menu to reach the [Block Storage Volumes](#) within VPC Infrastructure (IBM Cloud pull-down menu > VPC Infrastructure > Block Storage Volumes).
3. From the list of pre-existing **Block Volumes**, identify the volumes attaching to a Delphix Engine and wait until the volume state becomes Available.
4. Note that the volumes being attached have **Attachment Type** set as a hyphen.
5. Select the right side of the volume row menu, then select **Attach to Instance**.
6. In the **Attach Virtual Server Instance** modal window, select the virtual server instance (Delphix Engine) from the drop-down menu.
7. Click on the blue **Attach Volume** button.
8. Repeat Steps 3-7 until all associated **Block Volume** resources have been attached to the Delphix Engine instance.

Configuring the Delphix Engine

1. Connect to the running Delphix Engine instance with a web browser. Use the IP address or DNS name noted in the Instance Description. Upon successful connection, the browser will display a login prompt to enter the **Delphix Setup Page**.

2. Refer to the standard [product deployment instructions](#) to complete the Delphix deployment.

Next Steps

Congratulations! The Delphix Virtualization Engine should be successfully deployed in IBM Cloud.

Use Delphix documentation to learn how to:

- [configure the database source](#)
- [configure target environments](#)
- [create virtual databases \(VDBs\)](#)

Related Topics

- [Deployment for IBM Cloud](#)

Hotfix Information

Overview

If there are any hotfixes installed for the Delphix Engine, the list can be accessed via the system administrator CLI with the following commands.

```
engine> cd system
engine system> ls
....
hotfixes: [HF-111]
....
```

Configuration

 Unknown macro: 'navbox'

Registration Management

Registering your Delphix Engines is a prerequisite for serviceability from the Delphix Support team. Registration allows Delphix Support to access the engine and properly diagnose and identify any issues during support cases. It is important to manage your registration regularly in order to ensure the security of Support access, as described in the section below.

 Registering a Delphix Engine with support is different from registering it with Data Control Tower, *formerly* Central Management. For more information refer to [Data Control Tower](#).

The registration code contains an encrypted key that only Delphix can decrypt, which is unique for each engine. Delphix uses this key to generate one-time authentication codes that authorized Support personnel can use to log into the engine during support sessions.

Retrieving the Delphix Engine Registration Code

To enable Delphix Support, you may perform registration either during the initial setup of the Delphix Engine. If you want to register later, you may retrieve the registration code from the system setup portal.

1. You can retrieve the Delphix Engine Registration Code during the Initial Setup or later through the Delphix Setup application after logging in as a system administrator.
2. In the Registration panel, click View.
3. The Registration Code is displayed in the bottom half of the Registration window.
4. If your local workstation is connected to the external Internet, you can auto-register the Delphix Engine:
 1. Enter your Delphix Support Username and Support Password.
 2. Click Register.
 - If external connectivity is not immediately available, you must register manually.
 1. Copy the Delphix Engine registration code
 2. Log in with your support credentials at <https://register.delphix.com>.
 3. Paste the Registration Code and click Register.

Following registration, you will receive an email confirming the registration of your Delphix Engine

Regenerating the Delphix Engine Registration Code

Delphix recommends that you regenerate the registration code every six months to rotate the secret key in order to maximize the Support Security of the Delphix Engine.

Procedure

1. Log into the CLI (command-line interface) of the Delphix Engine with the sysadmin credentials.
2. Type `/registration/regenerate` and hit enter.
3. Type `commit` and hit enter. After a few seconds, the new code will be displayed.
4. Re-register the engine with this new code.

Failing to re-register the Delphix Engine after regenerating the registration code may prevent Support personnel from accessing the engine. In such a case, a support session cannot begin until the engine has been re-registered with the new registration code.

Related Topics

- [Configuration](#)
- [Installation and Initial System Configurations](#)

User and Authentication Management

There are two main topics in this section:

- Managing users and groups who will access and manage Delphix.
- Configuring various authentication mechanisms to access both Delphix and connected environments and datasets.

This section covers the following topics:

- [Users and Groups](#)
- [Authentication Mechanisms](#)

Users and Groups

User Types and User Management

There are three user types in the Delphix user model: the system administrator, the Delphix user, and the Self-service user.

i Usernames must start with a letter and contain only alphanumeric characters, hyphens, underscores, and/or periods.

System Administrators

System administrator users are responsible for managing the Delphix Engine itself, but not the objects (Environments, dSources, VDBs) within the server. For example, a system administrator is responsible for setting the time on the Delphix Engine and its network address, restarting it, creating new system administrator users (but not Delphix users), and other similar tasks.

A user called sysadmin is the default system, administrator user. While this user can be suspended, it may not be deleted. When the Delphix Management application first launches, this user can log in using the username sysadmin and password sysadmin.

To create or modify system administrators, first, log in to Delphix Setup and navigate to the Users section of the homepage. Here, you can:

- **Add new** system administrators with the plus sign
- **Change** system administrator passwords with the pencil icon
- **Delete** system administrators with the trashcan icon
- **Suspend** system administrators with the pause button
- **Reinstate** system administrators with the play button

Delphix Users

Delphix users are responsible for managing the environments and datasets within Delphix, such as dSources, virtual databases (VDBs), users, groups, and related policies and resources.

A Delphix user can be marked as a Delphix Administrator. Delphix Administrators have three special privileges:

- They can manage other Delphix users
- They implicitly have Owner privileges for all Delphix objects
- They can create new groups and new environments

The default Delphix user provided with a Delphix Engine is a Delphix Administrator and is called admin. Like the sysadmin user, the admin cannot be deleted. When the Delphix Management application launches, the admin user can log in using the password specified during the initial setup when Delphix was first launched.

Only these two users require password-based authentication. Also, other users may use other mechanisms such as LDAP or Kerberos, as described in [Configuring and Managing Kerberos](#) and [Configuring and using LDAP with the Delphix Engine](#).

Self-service Users

Delphix Self-Service has two types of users: the admin user and the data user.

Admin users have full access to all report data and can configure Delphix Self-Service, additionally, they can:

- Use the Delphix Engine to add/delete users

- Change tunable settings
- Add/delete tags
- Create and assign data templates and containers

Data users have access to production data provided in a data container. The data container provides these users with a playground in which to work with data using the Self-Service Toolbar.

For more information on Self-service users, visit our [Self-service documentation](#).

User Privileges for Delphix Objects

The user roles on Delphix objects consist of four types, which the Engine Admin user assigns: Provisioner, Owner, Data Operator, and Reader. These privileges apply both to objects, such as dSources and Virtual Databases (VDBs), and to groups, which are containers that hold those objects.

The Delphix Administrator user can assign privileges to groups, dSources, and VDBs. Privileges are inherited, meaning that privileges assigned to a group are effective for the dSources and VDBs contained in that group.

If a user does not have a privilege in relation to an object or group, then he or she has no visibility into that object or group.

Roles and Privileges for Delphix Objects

Role	Object Privileges	Group Privileges
Owner	<ul style="list-style-type: none"> • Can provision VDBs from owned dSources and VDBs • Can perform Virtual to Physical (V2P) from owned dSources and VDBs • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption • Can refresh or rollback VDBs • Can snapshot dSources and VDBs • Can start, stop, or re-start VDBs 	<ul style="list-style-type: none"> • Can provision VDBs from all dSources and VDBs in the group • Can refresh or rollback all VDBs in the group • Can snapshot all dSources and VDBs in the group • Can perform Virtual to Physical (V2P) from owned dSources and VDBs • Can view Templates for policies. • Can not create, edit, or delete a policy template from the policy page. • Can assign Owner privileges for dSources and VDBs • Can access the same statistics as a Provisioner, Data Operator, or Reader • Can start, stop, or re-start VDBs
Provisioner	<ul style="list-style-type: none"> • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption • Can provision VDBs from dSources and VDBs 	<ul style="list-style-type: none"> • Can access statistics on all dSources, VDBs, or snapshots in the group such as usage, history, and space consumption • Can provision VDBs from all dSources and VDBs in the group

Data Operator	<ul style="list-style-type: none"> • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption • Can refresh or rollback VDBs • Can snapshot dSources and VDBs 	<ul style="list-style-type: none"> • Can access statistics on all dSources, VDBs, or snapshots in the group such as usage, history, and space consumption • Can refresh or rollback all VDBs in the group • Can snapshot all dSources and VDBs in the group
Reader	<ul style="list-style-type: none"> • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption 	<ul style="list-style-type: none"> • Can access statistics on all dSources, VDBs, or snapshots in the group such as usage, history, and space consumption
Self-Service Only	<ul style="list-style-type: none"> • In the Delphix Self Service UI this user can: • Refresh • Restore • Bookmark • Reset • Branch • Stop/Activate • Share 	

Managing Groups

Creating groups helps you manage policies and privileges over objects within that group. When privileges are created for users at the group level, those privileges apply to all objects of that type within the group. When new objects are created or added to the group, the policies and privileges you have created at the group level will be applied to them.

- [Assigning Group and Object Ownership](#)
- [Adding and Deleting Groups](#)

Authentication Mechanisms

Delphix supports a variety of authentication mechanisms to connect to several different interfaces and systems. For example, you can connect via the UI using the default users described above, or you can connect to the CLI using an API token.

There are three categories of authentication related to Delphix: the Delphix UI, the Delphix CLI/API, and external systems such as Kerberos access to connected source and target hosts. Below are detailed pages related to each of these three sections:

- UI Authentication:
 - Data Control Tower, formerly Central Management
 - Username and password
 - LDAP: Directory-based authentication to Delphix engines rather than the default local access
 - Single Sign-on: Integration and support for identity providers to authenticate users on a per engine basis using SAML2-SSO.

- CLI Authentication:
 - Username and password
 - Auto-authentication via SSH keys: to automatically sign in to the Delphix CLI without requiring user-input credentials
- API Authentication
 - Username and password
 - API Tokens (for Delphix Engines registered with [Data Control Tower](#))
 - OAuth2 JSON Web Tokens
- External systems:
 - Username and password
 - SSH keys
 - Kerberos: Authentication for environments and data sources using Kerberos

Kerberos Support

Kerberos support is for access to connected environments, rather than the Delphix engine itself. This is an advanced topic and will require a solid understanding of Delphix concepts and architecture.

Related Topics

- [User and Authentication Management](#)
- [Users and Groups](#)
- [Assigning Group and Object Ownership](#)
- [Adding and Deleting Groups](#)

Assigning Group and Object Ownership

This topic describes how to assign group and object ownership to users in the Delphix Domain.

1. Log into the **Delphix Management** application.
2. Select **Manage > Users**.
3. For an existing user, select a user then select the edit icon.
4. Click **Next**.
5. In the **Privileges** tab assign **Owner** or **Provisioner** rights for groups or objects within groups. You do not have to assign a specific owner or auditor right for each object.
6. Click Next when finished.
7. For new users, refer to [Users and Groups](#). When you click **Submit**, the User Profile manager will reload, and then you can follow steps 4 - 6 to assign privileges.

Related Topics

- [Users and Groups](#)
- [Adding and Deleting Groups](#)

Adding and Deleting Groups

This topic describes how to add and delete groups within the Delphix Domain.

 User terminology for Delphix Admin has been changed to Engine Admin. The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Adding a Group

1. Log into the **Delphix Management** application.
2. From the **Manage** menu, select **Datasets**.
3. Select the **plus** icon and then select **Add Dataset Group**.
4. Enter a **Group Name** and an optional description.
5. Click **Add**.

At Least One Group Must Exist

At least one group must always exist on the Delphix Engine in order to link a dSource. If you delete the last group, you will need to create a new group in order to create a dSource.

Deleting a Group

1. Log into the **Delphix Management** application as a user with **Engine Admin** privileges or group **OWNER** privileges for the target group.
2. From the **Datasets** panel, select the target group.
3. Click the **Trash Can** icon.
4. Click **Delete**.

Deleting Groups Containing Objects

A group cannot be deleted if it contains VDBs or dSources. All databases within a group must be deleted prior to deleting the group.


Related Topics

- [Users and Groups](#)
- [Assigning Group and Object Ownership](#)

Managing System Administrators

Adding a New System Administrator

1. Launch the **Setup** application.
2. From the **Dashboard** select the **+** icon located in the **Users** card next to the filter field.



The screenshot shows a web interface for managing users. At the top, there is a header labeled "Users" with an information icon. Below the header, there are four action icons: a plus sign, a right-pointing triangle, a pencil, and a trash can. The main content area contains a table with two columns: "Username" and "Email". The table has one row with the values "sysadmin" and "manisha.gupta@delphix.com".

Username	Email
sysadmin	manisha.gupta@delphix.com

3. Enter the required information.

Add User ✕

Authentication Type

Local ▼

CREDENTIALS

Username

jsmith

Password

.....

Confirm Password

...

Email Address

USER DETAILS

First Name

Last Name

Phone (Work)

Phone (Cell)


Phone (Home)

Cancel Save


4. Click **Save**.

Deleting and Suspending System Administrators


1. Launch the **Delphix Setup** application and log in using the **sysadmin** (or another system administrator) credentials.
2. In the **User** panel, click the user you want to suspend or delete.
3. Suspend the user by clicking the **pause** icon ().



4. Once the user has been paused you can delete the user by clicking the **trash can**.



5. From the Delete User dialog select **Delete**.



Suspending the sysadmin User

The sysadmin user is a required user for the Delphix Engine. This user cannot be deleted but can be suspended. Suspending the sysadmin user prevents that user from being able to log into **Delphix Setup** or to the console via **ssh**.

Reinstating System Administrators

1. Launch the **Delphix Setup** application and login using system administrator credentials.
2. In the **User** panel, click on the name of the user you want to reinstate.
3. Reinststate the user by clicking the **enable user** icon (▶).



Changing System Administrator Passwords

1. Launch the **Delphix Setup** application and log in using **sysadmin** level credentials.
2. In the **User** panel, click the user whose password you want to change.



3. Select the **Edit** icon.

4. Enter the new password fields.

Edit User ✕

CREDENTIALS

Username
jsmith

Password
.....

Confirm Password
.....

Email Address

USER DETAILS

First Name

Last Name

Phone (Work)

Phone (Cell)

Phone (Home)

5. Click **Save**.

Adjust Session Timeout (Command-Line Only)

Configuring session timeout for System Setup users is not available in the GUI. As such, the only method is via CLI using the commands below.

```
Delphix> user
Delphix user> select sysadmin
Delphix user 'sysadmin'> ls
Properties
  type: User
  name: sysadmin
  apiUser: true
  authenticationType: NATIVE
  emailAddress: user.name@domain.com
  enabled: true
  firstName: (unset)
  homePhoneNumber: (unset)
  isDefault: true
  lastName: (unset)
  locale: en-US
  mobilePhoneNumber: (unset)
  passwordUpdateRequest: NONE
  principal: sysadmin
  publicKey: (empty)
  reference: USER-1
  sessionTimeout: 30min
  userType: SYSTEM
  workPhoneNumber: (unset)

Operations
delete
update
disable
enable
updateCredential
Delphix user 'sysadmin'> update
Delphix user 'sysadmin' update *> set sessionTimeout=180
Delphix user 'sysadmin' update *> commit
```

Related Topics

- [User and Authentication Management](#)

Managing Delphix Users

This section describes how to manage users. Here, you can learn how to:

- [Add Users](#)
- [Edit, Delete and Suspend Users](#)
- [Manage Profile Information](#)
- [Delphix User Account Lockouts](#)

Adding Users

Prerequisites

If you intend to validate user logins using LDAP authentication, make sure a system administrator has configured LDAP.

Procedure

1. Launch the **Delphix Management** application.
2. Click **Manage**.
3. Select **Users**.
4. Click plus icon to **Add User**.
5. Enter the mandatory fields **Username**, **Email Address**, and **New Password** for the new user.

Rules for creating a username

Your username:

- must be between 1 to 256 characters
- can be just letters, just numbers, or just any of the following special characters (., -, ,, @) or a combination of all of these. For example, a username could be just "@".
- can start with any of the above-listed characters and is case-sensitive.

Your password has no restrictions.

6. Select the **User Type**.
7. Click **Next**.
8. In the **Privileges** tab enter the privileges for the user.
9. Click **Next** and review the summary.
10. Click **Submit**.

Assigning Owner and Provisioner Privileges

Assigning owner privileges at the group level conveys ownership privileges over all objects in that group.

Click the **expand** icon next to each group name to see all objects in that group.

You can also assign ownership privileges only for specific objects in a group. You do not have to assign owner or auditor privileges for all Delphix objects, only those for which you want to grant the user-specific access.

Editing, Deleting, and Suspending Users

The delphix_admin user

The user named **delphix_admin** cannot be deleted since this is a user created by the Delphix Engine. However, you can suspend it.

 When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

1. Launch the **Delphix Management** application.
2. Select **Manage > Users**.
3. Click the user's name to open the user's profile panel.
4. Click the **disable** icon to disable the user.
5. Click the **trash can** icon to delete the user.

 Deleting a user cannot be undone.

Managing Individual Profile Information

1. After logging in, click your name in the menu bar.
2. Click **Profile**.
3. Edit profile information as necessary.
4. Select options for the event level that will trigger a notification email.
5. Select a time period for **Session Timeout**.
6. Click **Password** to edit your password.
7. Click **OK** when finished.
8. Click **Privileges** to see your privileges (Auditor or Owner) for Delphix objects.

Delphix User Account Lockouts

User Account Lockouts

This feature applies to all kinds of users -- Delphix and LDAP. It also applies to usernames that do not correspond to any user in the system. A user who enters a wrong password three times in a row is "locked out" (i.e., unable to continue attempting to log in) for an initial period of 30 seconds. After three more bad login attempts, the user must wait 60 seconds, then 90 seconds, and so on.

Troubleshooting a User Account Lockout

The initial wait time for any future lockouts is reset to 30 seconds when the user successfully logs in or when an administrator resets the user's password. When an administrator resets a locked-out user's password, the user can immediately attempt to log in.

Related Topics

- [User and Authentication Management](#)

Managing Individual Profile Information

This topic describes how individual users can manage personal settings such as personal information, passwords, event notifications, and session timeouts. It also describes how users can view their privileges for Delphix objects.

Procedure

1. After logging in, click your name in the menu bar.
2. Click **Profile**.
3. Edit profile information as necessary.
4. Select options for the event level that will trigger a notification email.
5. Select a time period for **Session Timeout**.
6. Click **Password** to edit your password.
7. Click **OK** when finished.
8. Click **Privileges** to see your privileges (Auditor or Owner) for Delphix objects.

Authentication Mechanisms

There are various authentication mechanisms that Delphix supports besides signing in via username and password to engines via the GUI, CLI or API. Below, you'll find documentation for the following capabilities:

- Auto-authentication via SSH keys: to automatically sign in to the Delphix CLI without requiring user-input credentials
- LDAP: Directory-based authentication to Delphix engines rather than the default local access.
- Single Sign-on: Integration and support for identity providers to authenticate users.
- Kerberos: Authentication for environments and data sources using Kerberos.
- OAuth2: Offers an alternative, password-less authentication method for API access to the Delphix Engine.

Kerberos Support

Kerberos support is for access to connected environments, rather than the Delphix Engine itself. This is an advanced topic and will require a solid understanding of Delphix concepts and architecture.

Related Topics

- [How to Setup Auto-authentication](#)
- [Configuring and using LDAP with the Delphix Engine](#)
- [Configuring Single Sign-on](#)
- [Configuring and Managing Kerberos](#)
- [Configuring OAuth2 Authentication for API Access](#)

How to Setup Auto-authentication

Generally, users need to enter a username and password when logging into the Delphix CLI. There are situations in which users may find entering a password cumbersome, or manual password entry may not be possible. These situations can be alleviated by setting up auto-authentication for the Delphix CLI.

There are two basic steps:

1. Generate a public and Private RSA key pair.
2. Register the public key with the specific Delphix Engine user.

There are two methods available:

- PuTTY
- OpenSSH with OpenSSL

If the examples provided do not work for you, you may need to consult your SSH documentation, we can only provide support for the Delphix Engine side of the connection. In both examples we grant password less login to the **sysadmin** user to host **Delphix5010**.

Using PuTTY

You will need both **putty.exe** and **puttygen.exe** for this.

Launch **puttygen.exe**

Set the **Type of key to generate** to *SSH-2 RSA* and the **Number of bites in generated key** to a suitable value such as *2048*. Click **Generate**

Once it has generated the key pair, leave the password fields blank and save the public and private keys to file.

Add the full contents of the public key to each Delphix Engine user you want to allow automatic login for.

```
Delphix5010> user
Delphix5010 user> select sysadmin
Delphix5010 user 'sysadmin'> update
Delphix5010 user 'sysadmin' update> set publicKey="ssh-rsa
AAAAB3NzaC1yc2EAAAABJQAAAQEAjdQYr1WU6UPr6FZqyt3eKNJEkAe8IdKQ8hcuBwa3HvRVmUuv0Lykm5AYQl
IW0B33aWusr0o+2FVTzt3/6G1llCf7wfhCShlJsYgwgMHeEGjixK5tacFCD8r+8dAlaXlv8u0lddK0A2LPXbCC
CIRL7IyVEnlSbUFY8s+E/
2R3owy5XSbLJLE1e15m1lQP0yUuQddAh25ruWR+1HHSaWG3p+wof0h6l7czkEcq7fPjtAZvivX90e8Ggt6JQ8b
v6td7aJW0bU2Y9YY0HLLHot7NQ4AT/
0tXSRKAG8sIdL7tY9hbHMNHRftCLzfn7mL+Qk8TjUYni3JGB4Vyi0bmkj6nHQ== rsa-key-20160309"
Delphix5010 user 'sysadmin' update> commit
```

1. In Putty, create a profile that uses the private key. In the PuTTY Connection settings set **SSH > Auth > Private key file for authentication** to the private key file you just generated
2. Next, still in the PuTTY Connection settings set **Data > Auto-login username** to `sysadmin@SYSTEM`.
3. Test the connection by setting the connection hostname as you normally would for PuTTY and click **Open**.

```
Connection > SSH . Auth
```

Using OpenSSH with Open SSL

Generally, OpenSSH will already have default public and private keys that can be used, if not (or the default keys are password locked) you can create them this way. OpenSSL is required but OpenSSH will take care of the background OpenSSL stuff for you.

1. Create your RSA key pair to a bit length suitable for your security needs (2048 is commonly required for recent security audits)

```
$ ssh-keygen -b 2048 -t rsa -P '' -f /etc/ssh/ssh_host_rsa_key
```

Results in a matching public file called `/etc/ssh/ssh_host_rsa_key.pub`

If you want to create different key pairs, just specify a different file path.

2. Add the full contents of the public key (`/etc/ssh/ssh_host_rsa_key.pub` in this example) to each Delphix Engine user you want to allow automatic login for.

```
Delphix5010> user
Delphix5010 user> select sysadmin
Delphix5010 user 'sysadmin'> update
Delphix5010 user 'sysadmin' update> set publicKey="ssh-rsa
AAAAB3NzaC1yc2EAAAABJQAAAQEAjdQYr1WU6UPr6FZqyt3eKNJEkAe8IdKQ8hcuBwa3HvRVmUuv0Lykm5AYQl
IW0B33aWusr0o+2FVTzt3/6G1lLCf7wfhCShlJsYgwgMHeEGjixK5tacFCD8r+8dALaXlv8u0lddK0A2LPXbCC
CIRL7IyVEnlSbUFY8s+E/
2R3owy5XSbLJLE1e15m1lQP0yUuQddAh25ruWR+1HHSaWG3p+wof0h6l7czkEcq7fPjtAZvivX90e8Ggt6JQ8b
v6td7aJW0bU2Y9YY0HLLHot7NQ4AT/
0tXSRKAG8sIdL7tY9hbHMNRftCLzfn7mL+Qk8TjUYni3JGB4Vyi0bmkj6nHQ== rsa-key-20160309"
Delphix5010 user 'sysadmin' update> commit
```

3. Test no password login on the command line from your client.

```
$ ssh -i /etc/ssh/ssh_host_rsa_key sysadmin@Delphix5010
```

Related Topics

- [Configuration](#)
- [SSH Implementation](#)

Configuring and using LDAP with the Delphix Engine

Using LDAP with the Delphix Engine requires the following:

- configure the Delphix Engine to use LDAP
- add LDAP users in the Delphix Management application

Configuring LDAP on the Delphix Engine

1. From the Delphix Setup application configure the LDAP server with the Delphix Engine by selecting Modify in the Authentication section.



2. Enter the information about the LDAP Authentication Server.

This must be an LDAP server that is configured for authentication. This information should come from the LDAP admin who runs the server. As a general rule only use simple auth. If using SSL/TLS typically use port 636 and import a certificate. If not using SSL/TLS, use port 389 and you will not need a certificate. If the remote LDAP server has disabled anonymous access and the user is trying to use SSL/TLS, the user will be unable to import the certificate. If this occurs file a support case so that Delphix Support can help manually upload the certificate.

Note

Import Server certificate option may import more than one certificate. It is recommended to import the CA certificate in the TrustStore and then click the Test Connection option to validate. For more information on adding a certificate in the TrustStore, refer to [TrustStore Settings](#).

Test Connect will issue an anonymous login request to the LDAP server. If the LDAP server has disabled anonymous access the test will fail. Test the server by adding a valid LDAP user and try logging in.

Authentication ✕

Use LDAP

LDAP Server

Port

636

Protect LDAP traffic with SSL/TLS

Authentication

SIMPLE ▼


Import Server Certificate

Test Connection


3. After updating the information and clicking the **Save** button, the *Authentication Service* section should reflect the proper information.

Create a new LDAP User account

1. Login to the Delphix Management application and go to **Manage > Users** to add a new user.



2. In the **Users** screen, click the **Add User** icon (**+**) and choose **LDAP** as the *Authentication Type*.



3. Fill out the data fields and decide if the user will be a *Delphix Admin*. For more info on the *Delphix Admin* setting please see [Managing System Administrators](#). When adding the principal, it is mandatory to specify the entire DN of the user to be added.

Each entry in an LDAP tree has a unique identifier; its Distinguished Name (DN). This consists of its Relative Distinguished Name (RDN), constructed from some attribute(s) in the entry, follow the parent entry's DN. Think of the DN as the full file path and the RDN as its relative filename in its parent folder (e.g. if `/foo/bar/myfile.txt` is the DN then `myfile.txt` would be the RDN). Some users prefer to use the term fully qualified DN to emphasize that a proper DN should include all of the components.

Example of LDAP Tree in which the base is:

```
dc=example,dc=com
```

and people are stored in a People subtree with RDN:

ou=people

and each individual is keyed by the cn (common name) attribute.

An example DN in this case would be: cn=Tony,ou=people,dc=example,dc=com

When adding an LDAP user you will be asked for the following information:

- Principal - which is the DN from above
- email address
- username - used to login into Delphix

A password is no longer required because it will authenticate against the password already stored in the LDAP entry, which is presumably known to the individual already. It is probably best if someone familiar with the LDAP tree and using it for authentication were involved at least initially to help understand how to describe the fully qualified DN for users.

Using Microsoft AD as an LDAP server

Using Microsoft AD as a modified LDAP server is also possible. Microsoft AD allows some shortcuts in the specification of the DN when binding.

Examples:

- <domain>\<user logon name>
- <user logon name>@<domain>.com

As with generic LDAP, it is probably best if someone familiar with using the AD LDAP instance for authentication was involved.

User

User Type

Standard User

Authentication Type

LDAP

CREDENTIALS

Principal

delphix\jwatson

Username

jwatson

Email Address

john.watson@delphix.com

When users log in, they will enter the username as chosen above, and the password that matches the principal entered above.

Related Topics

- [Configuration](#)
- [User and Authentication Management](#)
- [Authentication Mechanisms](#)

Configuring Single Sign-on

Overview

This article provides instructions on how to set up Single Sign-on (SSO) on the 5.3.3+ Delphix Engine. Delphix Engines (Masking and Virtualization) versions 5.3.3+ support authentication via the SAML 2.0 standard (SP initiated and IdP Initiated). SLO (Single log-out) is not supported. This means that logging out of a Delphix Engine will not terminate sessions on other Delphix Engines, nor will it terminate the IDP session.

Identity Provider Configuration

The steps to configure an Identity Provider (IdP) are specific to each IdP product (e.g. Okta, OneLogin, PingFederate). The terminology may vary, but one SAML 2.0 application (or SP connection) will need to be created for each Delphix Engine. The engine does not expose a metadata document. The following attributes must be entered:

- ACS URL (Assertion Consumer Service URL): `http(s)://<delphix-engine>/sso/response`
 - Delphix strongly recommends that HTTPS is used instead of HTTP for all UI and API communication with the Delphix Engine. For HTTPS or the automatic HTTP to HTTPS redirect, use the **https://** scheme in the ACS URL, otherwise use the **http://** scheme. Refer to the [CLI Cookbook: Changing HTTP and HTTPS Web Connections](#) article on how to set up HTTPS.
- SAML Bindings: Delphix Engines support the POST and Redirect bindings.
- Audience Restriction (SP entity ID, Partner's Entity ID): The audience restriction must be set to the entity id configured in the Delphix Server via the Delphix Setup (see below). The default value is **https://<Delphix Server ID>** where <Delphix Server ID> is a 36-character hexadecimal string of the form `xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx`. See [Determining the Delphix Server ID and Host Name](#) for more on the Delphix Server ID.
 - If the Delphix Engine does not exist or is unreachable, enter a temporary value (such as `delphix-sp-id`) to later be replaced by the actual Delphix Server ID.
- Signature policy: The Delphix Engine does not sign authentication requests; it requires that either the responses, assertions, or both are signed.
- Name ID: The SAML NameID attribute must be set to the email address of the user. See the *User Management When SSO is Enabled* section below for more information.

The SP initiated flow must be enabled in the IDP.

New Engine Configuration

Follow this procedure when installing a new Delphix Engine.

1. Connect to the Delphix Engine at `http://<DelphixEngine>/login/index.html#serverSetup`. The **Delphix Setup** application will launch when connecting to the server.
2. Enter the **sysadmin** login credentials; this account has a default username of **sysadmin** and password of **sysadmin**.
3. In the Authentication step of the Delphix Setup wizard, check the **use SAML/SSO** box and enter the required information:
 - a. The entity id is a unique identifier of the Delphix Engine for Single Sign-On providers. The default value is **https://<Delphix Server ID>**, where <Delphix Server ID> is a 36-character hexadecimal string of the form `xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx`. This is just an identifier; there is no resource at that URL. This value can be changed to any string, but note that some identify providers require this to have a URL format.
 - b. The IdP metadata is an XML document that must be exported from the application created in the IdP (see [Identity Provider Configuration](#)).

- c. Optional **Advanced** settings include the **response skew time** which is the maximum time difference allowed between a SAML response and the engine's current time, in seconds. If not set, it defaults to 120 seconds. The **maximum age of IdP authentication** indicates how far in the past to accept authentications to the identity provider, in seconds. If not set, it defaults to 86,400 seconds (one day).
4. Complete the remaining setup steps as usual.

Existing Engine Configuration

Follow this procedure to enable SSO on an already configured engine.


1. Connect to the Delphix Engine at `http://<DelphixEngine>/login/index.html#serverSetup`. The **Delphix Setup** application will launch once connected to the server.
2. Enter the **sysadmin** login credentials.
3. In the **Authentication** tile, select **Modify**.



4. In the Authentication dialog, check the use SAML/SSO box, then enter the entity id (if not using the default), IdP metadata, and the optional advanced time settings (described in [New Engine Configuration](#)).

⚠ Engine Type

If this is an upgraded engine, make sure the engine type is set from the banner of the ServerSetup application dashboard.



User Management When SSO is Enabled

Access to the Delphix Setup application is not affected by the use of SSO, it only affects access to the Delphix Management Application and Masking Application. When SSO is enabled, authentication to the Delphix Management Application or Masking Application UIs are performed via SAML/SSO instead of a combination of username and password. Non-administrators can no longer change their email address.

An administrator must create a Delphix user for each user to whom access via SSO must be granted. The Delphix user can be used to assign roles and permissions. The email address of the Delphix user is used to match users authenticated via SAML/SSO and must be set to the exact value defined in the IdP. This same value is used in the NameID attribute of the SAML response. Supported NameID formats

are `urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified` and `urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress`.

If multiple Delphix users share an email address, all Delphix users will have access to the SAML/SSO session.

API Access

Neither the API access for use in scripts and integrations nor the Virtualization CLI access requires SSO. Instead, username and password (optionally with LDAP integration) authentication must be used for API or CLI access. When SSO is enabled on a Delphix Engine, new users by default will have no API access and no password. Administrators can enable API access for any user through the User Management or Masking UIs. The user's password or LDAP credentials are used for API authentication.

Users created before enabling SSO will maintain their API access enabled. Users with API access but no email address are useful for scripts or integration via the API - they cannot be used to log in via the UI. When SSO is enabled, only administrators can change or set email addresses.

A user with API access may also log in via SAML into an SSO-enabled engine through the UI when they have an email set.

Troubleshooting

If authentication via SSO fails with a message stating that the issue time is either too old or in the future, the error is due to the time on the Delphix Engine not being in sync with the time of the Identity provider server. If the time on the Delphix Engine is not correct, correct the time settings manually or configure NTP. Alternatively, update the response skew time parameter (see [Existing Engine Configuration](#)).

If authentication via SSO fails with a message stating that the authentication to the identity provider is too old, re-authenticate with the identity provider or adjust the maximum age parameter (see [Existing Engine Configuration](#)).

Related Topics

- [Configuration](#)
- [User and Authentication Management](#)

Configuring and Managing Kerberos

i **Version 6.0.7.0 or Later Recommended for Kerberos**

Any Delphix Engine intending to leverage Kerberos credentials should be running version 6.0.7.0 or later. Versions 6.0.0.0-6.0.6.1 may encounter issues in authentication ticket renewal, causing Environment and Dataset job failures. More information can be found in this Delphix [Knowledge Base](#) article.

This section covers the following topics :

- [Delphix Kerberos Implementation](#)
- [SSH Implementation](#)
- [Kerberos Requirements](#)
- [Configuring Kerberos](#)

Delphix Kerberos Implementation


Version 6.0.7.0 or Later Recommended for Kerberos

Any Delphix Engine intending to leverage Kerberos credentials should be running version 6.0.7.0 or later. Versions 6.0.0.0-6.0.6.1 may encounter issues in authentication ticket renewal, causing Environment and Dataset job failures. More information can be found in this Delphix [Knowledge Base](#) article.

Shared Infrastructure/Ticket Management

The Delphix Engine (DE) has a single Kerberos principal shared between all connections to the host (SSH, ASE JDBC, etc).

Overview of the Authentication Process



1. The client acquires a ticket from the Kerberos Domain Controller (KDC) (e.g. `kinit <principal>`) which it stores locally.
2. The client uses a ticket from KDC to authenticate with the target (e.g., ssh- or JDBC authentication using `gssapi` to pass the cached ticket acquired in step 1).

Kerberos Master/Replica KDCs

Kerberos supports a master/replica system with multiple KDCs running on different hosts. This is used for High Availability (HA) or to provide faster service via a local node in dispersed network environments. Delphix supports a list of KDCs for the Kerberos realm to which it has been joined.

Delphix Infrastructure to Support the Authentication Process

Kerberized Environment User

Delphix has introduced a `KerberosCredential` type that indicates the global Kerberos principal to be used for authentication, rather than user-specific credentials.

Keytab Based Authentication

It is possible to use `kinit` with a keytab file instead of password-based authentication to acquire tickets. This is similar in principle to passwordless SSH authentication and allows Delphix to function in the customer's environment without storing any passwords on the Delphix Engine. It does, however, put us at the mercy of the customer's keytab expiration policy.

The Delphix Engine creates a background thread that periodically checks the expiration of the cached Kerberos credentials. If the credentials have expired, it calls `kinit` using the keytab that was provided.

Keytab File Storage

Keytab file data is sent via a web service API as a Base64 encoded string. This is then decoded back to the binary file and persisted on local storage on the Delphix Engine with root user ownership permissions.

Default Behavior

Default Kerberos ticket refresh configuration:

- Delphix checks if the TGT-cached Kerberos ticket should be refreshed every hour.
- The TGT-cached ticket for the global Delphix principal will be refreshed if it expires in less than two hours. The default values can be changed by Delphix Support.


Related Topics

- [Configuring and Managing Kerberos](#)
- [Kerberos Requirements](#)

SSH Implementation

The management stack uses `sshj+gssapi` to pass already-generated Kerberos tickets to the Kerberized `sshd` on the source/target side if prompted to do so by the end-user passing a Kerberized environment user to existing wrapper functions.


⚠ The only thing changing from password-based or regular passwordless SSH authentication is the authentication step. Command execution remains unchanged.



SAP ASE, Oracle, and DB2 Connections

Delphix connects to SAP ASE, Oracle, and DB2 instances using the two listed mechanisms below. This example configuration uses an SAP ASE instance.

- via `isql` process
- via the `jConnect` JDBC driver



When connecting via isql Delphix uses the “-V” parameter rather than specifying a username/password. The “-V” option uses the Kerberos principal in the current user’s cached credentials file. Delphix relies on the end customer to configure this appropriately for their environment (for example, the cached credentials could be populated by a PAM module during login). Delphix also expects that the `KRB5CCNAME` is set appropriately or the credential cache is in the host default location.

When connecting via JDBC, Delphix uses additional connection

options: `REQUEST_KERBEROS_SESSION=true&SERVICE_PRINCIPAL_NAME=<ASE Instance SPN>`.

By default, the instance Service Principal Name (SPN) is identical to the instance name for authentication. Delphix allows the instance SPN to be manually set on a per-repository basis to allow for non-default values. The jConnect JDBC driver connects using the cached credentials that were obtained as described in the Shared infrastructure/ Ticket Management section.

For example, if the instance name is `ASE_INSTANCE_1` and has been configured to use `REALM.COM`, then the instance will attempt to authenticate with the KDC using `ASE_INSTANCE_1@REALM.COM`. However, this is configurable and can be specified either via an environment variable or a command-line option to the data server process. If an environment variable is used to configure the SPN, the instance must be manually discovered via web service APIs or the Delphix CLI.

Related Topics

- [Configuring and Managing Kerberos](#)
- [Configuration](#)

Kerberos Requirements

Prerequisites

Basic requirements prior to the configuration:

- MIT Kerberos 1.4.4 KDC
- Kerberos `REALM` name
- Global Kerberos principal name (specified without trailing `@REALM` name)
- Global Kerberos principal `keytab` data encoded as a base 64 string
- KDC hostnames and port numbers (one or more in priority list order)

Environment Requirements

i Version 6.0.7.0 or Later Recommended for Kerberos

Any Delphix Engine intending to leverage Kerberos credentials should be running version 6.0.7.0 or later. Versions 6.0.0.0-6.0.6.1 may encounter issues in authentication ticket renewal, causing Environment and Dataset job failures. More information can be found in this Delphix [Knowledge Base](#) article.

The following hosts and software versions are required:

- A source host with the following configuration:
 - A running SAP ASE, Oracle, or DB2 instance.
 - A database to link from and its corresponding full database dump.
 - The Delphix principal is able to access the SAP ASE instance and SSH onto the host as per our product documentation (see [Requirements for SAP ASE](#)).
 - The credential cache for the Delphix principal is populated and kept current. The environment variable `KRB5CCNAME` is set to the location for a credential cache. Login to the Adaptive Server via `" isql_r64 -V -R<spn> -S<servername> "` or `" isql_r -V -R<spn> -S<servername> "` or otherwise make sure that `" isql "` points to either `" isql_r64 or isql_r "` so that `" isql -V -R<spn> -S<servername> "` works.
- A staging host with the following configuration:
 - A running SAP ASE, Oracle, or DB2 instance with the same version as the source instance.
 - The Delphix principal is able to access the SAP ASE instance and SSH onto the host as per our product documentation (see [Requirements for SAP ASE](#)).
 - The credential cache for the Delphix principal is populated and kept current. The environment variable `KRB5CCNAME` is set to the location for a credential cache. Login to the Adaptive Server via `" isql_r64 -V -R<spn> -S<servername> "` or `" isql_r -V -R<spn> -S<servername> "` or otherwise make sure that `" isql "` points to either `" isql_r64 or isql_r "` so that `" isql -V -R<spn> -S<servername> "` works.
- A target host to create a VDB on. Configuration details:
 - A running SAP ASE, Oracle, or DB2 instance with the same version as the source instance.
 - The Delphix principal is able to access the SAP ASE instance and SSH onto the host as per our product documentation (see [Requirements for SAP ASE](#)).

- The credential cache for the Delphix principal is populated and kept current. The environment variable `KRB5CCNAME` is set to the location for a credential cache. Login to the Adaptive Server via "`isql_r64 -V -R<spn> -S<servername>`" or "`isql_r -V -R<spn> -S<servername>`" or otherwise make sure that "`isql`" points to either "`isql_r64` or `isql_r`" so that "`isql -V -R<spn> -S<servername>`" works.

Supported Databases and Kerberos Configurations

For detailed Kerberos support please refer to [Kerberos Support Matrix](#).

Delphix Engine 5.3.2.0 is the first generally available customer release to support Kerberos on a subset of supported OS' with SAP ASE, Oracle, and DB2 databases.

 With the Delphix Engine 5.3.2.0 release, many of the previously unsupported GUI functions now function when Kerberos is enabled. See [Configuring Kerberos via the UI](#) below for details, including limitations.

Related Topics

- [Configuring and Managing Kerberos](#)
- [Add a Kerberos Environment](#)

Configuring Kerberos

This section covers the following topics:

- [Configuring Kerberos during Engine Setup](#)
- [Add a Kerberos Environment](#)
- [Adding and Linking a dSource](#)

Configuring Kerberos during Engine Setup

Version 6.0.7.0 or Later Recommended for Kerberos

Any Delphix Engine intending to leverage Kerberos credentials should be running version 6.0.7.0 or later. Versions 6.0.0.0-6.0.6.1 may encounter issues in authentication ticket renewal, causing Environment and Dataset job failures. More information can be found in this Delphix [Knowledge Base](#) article.


A Network Authorization step has been added to the Delphix Engine setup wizard allowing users to use Kerberos authentication along with another authentication method.

1. From the Delphix Setup wizard Welcome screen select **Next**.
2. Complete the following tabs as you normally would.
3. In the **Network Authorization** tab select the **Use Kerberos authentication to communicate with the remote hosts** checkbox.
4. Complete the following fields:
 - Realm
 - Principal
 - Keytab

The screenshot shows the 'DELPHIX SETUP' wizard interface. The 'Virtualization Setup' section is active, and the 'Network Authorization' step is selected in the left-hand navigation pane. The main content area is titled 'Network Authorization' and 'KERBEROS CONFIGURATION'. A checkbox labeled 'Use Kerberos authentication to communicate with remote hosts' is checked. Below this, there is a section for 'Kerberos Key Distribution Center host(s)' with a table that currently has no rows. The table has columns for 'Hostname' and 'Port'. Below the table, there are three text input fields: 'Realm', 'Principal', and 'Keytab', each containing the placeholder text '(none defined)'. A 'Setup Help' link is visible in the top right corner.

5. Then select **Next**.


6. Confirm your selections and select **Submit**.



Editing Kerberos Configuration Settings

To edit your configuration settings after setup, login as a sysadmin user and complete the following steps:

1. On the Delphix SETUP Dashboard, from the **Network Authorization** panel, click **Modify**.



2. Select the check box under KERBEROS CONFIGURATION, select the hostname and edit the configuration.

Network Authorization

KERBEROS CONFIGURATION

Use Kerberos authentication to communicate with remote hosts

Kerberos Key Distribution Center host(s)

+ 🗑️

Hostname	Port
kerberos-02.delphix.com	88

Realm

Principal

Keytab

HOST CONNECTION AUTHENTICATION

When connecting to hosts, you can provide username-password pairs when setting up the connection, or you can utilize one or more Enterprise Password Vault systems by adding them to your engine setup.

Click the + to add a vault

3. Click **Save**.

Related Topics

- [Configuring and Managing Kerberos](#)
- [Delphix Kerberos Implementation](#)
- [Kerberos Requirements](#)

Add a Kerberos Environment

Version 6.0.7.0 or Later Recommended for Kerberos

Any Delphix Engine intending to leverage Kerberos credentials should be running version 6.0.7.0 or later. Versions 6.0.0.0-6.0.6.1 may encounter issues in authentication ticket renewal, causing Environment and Dataset job failures. More information can be found in this Delphix [Knowledge Base](#) article.

It is possible to create a Unix/Linux Standalone environment with Kerberos authentication.

1. Login to the **Delphix Management** application as an **admin**.
2. From the **Manage** menu, select **Environments**.
3. Then click on the plus icon to open a wizard to create a new environment.
4. Select your **Host OS** and **Server Type**, then select **Next**.
5. Under **Login Type**, select **Kerberos Authentication**.
6. If **Discover SAP ASE** is enabled, ASE DB Kerberos authentication will be available, select **Kerberos Authentication**.

The screenshot shows a configuration window with the following settings:

- Discover SAP ASE**: Enabled
- Login Type**:
 - Username and Password
 - Kerberos Authentication
- Principal**: sybase

7. Select **Submit**.

Changing the Environment User

1. Login to the **Delphix Management** application as an **admin**.
2. From the **Manage** menu, select **Environments**.
3. Select an Environment, the **Details** tab allows you to see Environment information.
4. On the grid of Environment Users, you can see existing users. Click the **plus icon**, to add a new user.
5. It is now possible to create a user with Kerberos Authentication. There can only be one Kerberos user per environment. The Principal is taken to the Kerberos Configuration.

It is possible to set the Kerberos user as Primary.

Changing SAP ASE DB User and DB Password

1. Login to the **Delphix Management** application as an **admin**.
2. From the **Manage** menu, select **Environments**.
3. Select an Environment, the **Details** tab allows you to see Environment information.

- Click the SAP ASE Information **pencil** icon.



- Click the edit icon and edit ASE DB User and ASE DB Password.
- Click the **checkmark** to save.

Updating Environment Notes

- Login to the **Delphix Management** application as an **admin**.
- From the **Manage** menu select **Environments**.
- Select an Environment, the **Details** tab allows you to see Environment information.
- Under **SAP ASE Information**, you can see **Notes**. To edit click the **pencil** icon.
- Click the **checkmark** to save.

Changing Host Address, SSH Port Number, and Toolkit Path

- Login to the **Delphix Management** application as an **admin**.
- From the **Manage** menu select **Environments**.
- Select an Environment, the **Details** tab allows you to see Environment information.
- Click the **pencil** icon located next to **Attributes** to edit the Host Address or SSH Port number or Toolkit Path.
- Select the **checkmark** to save.


Related Topics

- [Configuring and Managing Kerberos](#)
- [Kerberos Requirements](#)


Adding and Linking a dSource

Adding a dSource

1. Login to the **Delphix Management** application as an **admin**.
2. From the **Manage** menu select **Environments**.
3. Select an Environment and click the **Databases** tab. This tab provides details on all the available repositories.
4. Click on the **Add dSource** link.



5. Click on a **Data Source**, then select **Kerberos Authentication**.



6. Complete all remaining fields, and then click **Submit**.

Linking a dSource

1. Login to the **Delphix Management** application as an **admin**.
2. From the **Manage** menu select **Datasets**.
3. Select a **Dataset**.
4. From the Actions menu located on the top-right select **Link dSource**.
All information in the Link dSource screen is automatically filled in with the data you previously selected.

Related Topics

- [Configuring and Managing Kerberos](#)
- [Kerberos Requirements](#)

Configuring OAuth2 Authentication for API Access

Overview

This article provides instructions on how to set up Open Authorization 2.0 (OAuth2) on the 6.0.11 Delphix Engine. OAuth2 offers an alternative, password-less authentication method for API access to the Delphix Engine.

Delphix Engine (Masking and Virtualization) version 6.0.11.0 supports authentication using JSON Web Tokens (JWTs) issued by a known authorization server or identity provider (IdP). It is necessary for JWTs to contain a claim that can be used to associate an authentication request with a user that exists in the Delphix Engine. This article describes how to configure the Delphix Engine to validate tokens and associate token claims with Delphix Engine users.

 The authentication feature described in this section differs from the API token authentication feature supported by Delphix Engines registered with Data Control Tower (*formerly* Central Management). For more information on API token authentication, refer to [Data Control Tower](#).

Configuration Options

The following options for configuring OAuth2 for the Delphix Engine are available in the Delphix CLI (as the sysadmin user, under `service; oauth2`) as well as via the API endpoint `/resources/json/delphix/service/oauth2`.

Option	Description	Default (if applicable)
<code>audience</code>	Specifies the expected value of the audience claim (<code>aud</code>) of JWTs indicating that the tokens are intended for this particular Delphix engine.	<code>api://delphix</code>
<code>enabled</code>	Specifies whether the OAuth2 feature should be enabled.	<code>false</code>
<code>issuerURI</code>	(Required) Specifies the base location or identifier of the authorization server (IdP) which the Delphix Engine will use to validate incoming JWTs.	
<code>jwkSetURI</code>	Specifies the URI used to retrieve the JSON Web Key (JWK) set, if supported by the authorization server (IdP).	
<code>tokenSkewTime</code>	Specifies the maximum time difference (in seconds) allowed between the validity period of a JWT and the engine's current time.	<code>60</code>

Option	Description	Default (if applicable)
<code>userIdClaim</code>	Specifies which claim in a JWT should be used by the Delphix Engine (in conjunction with the <code>userMatchingFieldType</code> setting) to associate a token with a user configured on the Delphix Engine. The default (<code>sub</code>) corresponds to the subject claim of the token.	<code>sub</code>
<code>userMatchingFieldType</code>	Specifies which property of a Delphix Engine user will be used to match with the claim (specified in <code>userIdClaim</code>) of a JWT. The Delphix Engine can be configured to match users based on a user's <code>name</code> , <code>emailAddress</code> , or <code>principal</code> properties.	PRINCIPAL

These options may also be set in the Delphix Setup application. Please refer to the following procedures for configuring OAuth2 parameters for new or existing Delphix Engines.

OAuth2

Use OAuth2 access tokens

Issuer URI

`https://myissuer.com`

URI of the OAuth Authorization issuer.

Audience

`api://delphix`

The intended audience of the access tokens issued by the Authorization Server for Delphix Engine access.

User Identifying Claim

`sub`

The claim in a token that should be used to associate a JWT with a Delphix Engine user.

Show advanced ▼

New Engine Configuration

Follow this procedure to configure OAuth2 on a new Delphix Engine.


1. Connect to the Delphix Engine at `http://<DelphixEngine>/login/index.html#serverSetup` . The **Delphix Setup** application will launch when connecting to the server.
2. Enter the **sysadmin** login credentials; this account has a default username of **sysadmin** and password of **sysadmin**.

3. In the Authentication step of the Delphix Setup wizard, check the **Use OAuth2 access tokens** box and configure the OAuth2 settings as desired. (Refer to the table above for more information on these settings.)
4. Complete the remaining setup steps as usual.

Existing Engine Configuration

Follow this procedure to configure OAuth2 on an existing Delphix Engine.

1. Connect to the Delphix Engine at `http://<DelphixEngine>/login/index.html#serverSetup`. The **Delphix Setup** application will launch when connecting to the server.
2. Enter the **sysadmin** login credentials.
3. In the **Authentication** tile, select **Modify**.



4. In the Authentication dialog, check the **Use OAuth2 access tokens** box and configure the OAuth2 settings as desired. (Refer to the table above for more information on these settings.)

Example of API Access using OAuth2 Token

1. Obtain a JWT from the Authorization Server (IdP). (The details for this process will vary depending on the IdP vendor.) For the purposes of this example, the contents of the token are stored in the environment variable **t**.
2. Access the oauth2-login API endpoint of the Delphix Engine, providing the OAuth2 token. In this example the session information is stored in the file **cookies.txt** in the working directory. For Virtualization Engines the oauth2-login API endpoint is `/virtualization/api/oauth2-login`. For Masking Engines the oauth2-login API endpoint is `/masking/api/oauth2-login`.

Virtualization API Endpoint Example

```
h=<engine address>; curl -i -X POST $h/virtualization/api/oauth2-login -H
"Authorization: Bearer $t" -b cookies.txt -c cookies.txt -H 'Content-Type:
application/json' -d '{"type": "APISession", "version": {"type":
"APIVersion", "major": 1, "minor": 11, "micro": 11}}'
```

Masking API Endpoint Example

```
h=<engine address>; curl -i -X POST $h/masking/api/oauth2-login -H
"Authorization: Bearer $t" -H 'Content-Type: application/json' -b cookies.txt
-c cookies.txt
```

3. Refer to the saved cookies.txt file in subsequent **curl** invocations. The example below can be used to list the users configured on the Delphix Engine.

Virtualization API Example

```
curl -X GET -b cookies.txt -c cookies.txt -H 'Content-Type: application/json'
$h/resources/json/delphix/user
```

Masking API Example

```
# Store authorization code returned by /masking/api/oauth2-login in $m
curl -i -X GET -b cookies.txt -c cookies.txt -H 'Content-Type: application/
json' -H "Authorization: $m" $h/masking/api/v5.1.11/users
```

CLI Access using OAuth2 Token

When OAuth2 is enabled, CLI logins will prompt the user to supply either an OAuth2 access token or a password to authenticate:

CLI Access Prompt with OAuth2 Enabled

```
$ ssh myengine -l admin
Enter access token or press enter to provide a password:
```

To authenticate using an OAuth2 token, paste its contents when this prompt is shown. As with API authentication, the OAuth2 token must be current (not expired) and must contain a claim that can be associated with a valid Delphix Engine user, based on the **userIdClaim** and **userMatchingFieldType** values set on the Delphix Engine.

To bypass OAuth2 authentication and use password authentication, press **Enter** when this prompt is shown and a conventional password prompt will be displayed.

User Matching Policy

If the **userIdClaim** component of a JWT matches more than one Delphix Engine user (for example, if **userMatchingFieldType** is set to **EMAIL_ADDRESS**, and the same email address is associated with multiple Delphix Engine users), the oldest user account (by time of creation) will be authenticated.

 Suggestion

To ensure all users can be authenticated using OAuth2, make sure that the property specified in **userMatchingFieldType** is populated and unique for all Delphix Engine users.

HTTPS Proxy Configuration

If the OAuth2 identity provider cannot be directly reached by a Delphix Engine, a HTTPS proxy may be used. The sysadmin user can create or modify HTTPS proxy settings via the CLI/API endpoint `service; proxy`, or in the **Delphix Setup** application by selecting to modify the **Outbound Connectivity** tile, checking the **Configure web proxy** box, and entering the host, port, and (optionally) username and password of the proxy server.

Network and DNS Management

You may want to manage and configure certain network services, such as DNS, for Delphix. Here we specify general network and connectivity requirements, as well as detail how you can test network performance. [General Network and Connectivity Requirements](#).

This section covers the following topics:

- [General Network and Connectivity Requirements](#)
- [Network Performance Configuration Options](#)
- [Determining the Delphix Server ID and Host Name](#)
- [Configuring Multiple DNS Domain Names in DNS Search List](#)
- [How to Change the IP Address of the Delphix Engine](#)
- [How to Change the Hostname of the Delphix Engine](#)
- [How to Change the DNS Server of the Delphix Engine](#)
- [Configuring a Second Network Interface](#)

General Network and Connectivity Requirements

Overview

This topic covers the general network and connectivity requirements for the Delphix Engine, including connection requirements, port allocation, and firewall and Intrusion Detection System (IDS) considerations. For platform-specific network and connectivity requirements, see the relevant topics under the **Requirements** section for each platform.

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine

Protocol	Port Number	Use
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

Related Topics

- [Network and DNS Management](#)
- [Configuring Replication](#)

Network Performance Configuration Options

These topics describe configuration options to maximize network performance of a Delphix Engine deployment.

- [Optimal Network Architecture for the Delphix Engine](#)
- [Network Operations Using the Delphix Session Protocol](#)
- [Network Performance Test Tool Interface](#)
- [Working with Dataset Performance](#)
- [Network Performance Expectations and Troubleshooting](#)

Optimal Network Architecture for the Delphix Engine

This topic describes basic network performance considerations for the Delphix Engine.


Network Architecture and Latency

All VDB I/O operations are serviced over the network. Delphix uses NFS as the primary transport for Oracle VDBs, and iSCSI for MS SQL VDBs. The network architecture, latency, and capacity between the Delphix Engine and the target environment are key network components for improving the performance of a Delphix deployment. The latency between the Delphix Engine and the source environment is not relevant for the best performance of VDBs.


For optimal performance of VDBs, round-trip latency between the Delphix Engine and the target environment should be kept under 1 millisecond, and preferably in the range of 300 microseconds. If network latency exceeds 500 microseconds, the VDBs will not perform as well as a database connected to physical storage.

Latency can be introduced by having to route the network packets across multiple networks, or by the presence of routers, switches, and firewalls between the Delphix Engine and the target environment. Best practices to reduce network latency include:

- Keep the Delphix Engine on the same subnet as the target environment
- Reduce the number of hops between the Delphix Engine and the target environment
 - Reduce the number of switches in the network. Each switch can add 50 - 100 microseconds of latency to the network.
 - Reduce the number of routers in the network. Each router can add 500 - 1000 microseconds of latency in a network, and the round trip for an I/O operation could increase by as much as 1 - 2 milliseconds.
- There should be no firewalls between the Delphix Engine and the target environment.
- When linking the Delphix Engine to a source database across a WAN, consider the time needed for the initial link and load. It may be necessary to schedule the load operation as multiple steps across multiple days.



A Common WAN Deployment Architecture



Deployment of the Delphix Engine on Separate Sub-Nets

Network Throughput and Bandwidth

Network throughput measures the rate at which data can be sent continuously between two servers on a network. Network throughput is affected by network latency, but the dominant factor affecting throughput is the bandwidth of the network. As a point of comparison, consider the bandwidth available for three types of Ethernet networks:

Ethernet Type	Network Bandwidth
100Mb Ethernet (100Base-T)	~=10MB/sec
Gigabit Ethernet (GbE)	~=100MB/sec
10 Gigabit Ethernet (10GbE)	~=1GB/sec

Low network throughput can impact the Delphix Engine in a number of ways:

- Increasing the amount of time it takes to perform a SnapSync operation, both for initial load and subsequent regular snapshots
- Managing LogSync operations in a high change environment
- Poor VDB performance when an application is performing large sequential I/O operations, such as sequential table scans for reporting or business intelligence, or RMAN backups of the VDB.

Delphix Engine throughput must exceed the sum of the peak I/O loads of all VDBs. Delphix incorporates an I/O-Collector toolkit to collect I/O data from each production source database and pre-production server.

Best practices to improve network throughput include:

- Use 10 Gigabit Ethernet (10GbE)
- Use a dedicated storage network

If you are concerned about your network throughput, you can test it with the built-in CLI tool for network testing.

Related Topics

- [Network Performance Configuration Options](#)
- [Network Operations Using the Delphix Session Protocol](#)
- [CLI Cookbook: Network Performance](#)

Network Operations Using the Delphix Session Protocol

This topic describes how the Delphix Engine uses the Delphix Service Protocol (DSP) for network operations, and how this affects features such as replication, V2P, and SnapSync.

Overview

Delphix Session Protocol, or DSP, is a communication protocol that operates at the session and presentation layer in the Open Systems Interconnection (OSI) model.

Application Layer	application specific logic
Presentation Layer	data encoding, digest, compression, encryption
Session Layer	connection management, error recovery, security, remote operation
Transport Layer	end-to-end connection, message segmentation, sequencing, reliability, flow control
Network Layer	packet fragmentation, routing, logical addressing
Data Link Layer	physical addressing
Physical Layer	media, signal, binary transmission

DSP supports the request-reply pattern for communication between two networked peers. In contrast to the traditional remote procedure call (RPC) models, which focus exclusively on low-level details such as data encoding and wireframing, DSP implements a generic session layer that supports a number of advanced functionalities desired for network communication, including:

- Full-duplex remote operation execution and end-to-end cancellation support
- Advanced connectivity model with connection trunking and ordered delivery
- Fault resilience with automatic connection and session recovery, exactly-once semantics, and optional data digest
- High performance with concurrent execution, session flow control, optional data compression, and bandwidth throttling
- Built-in security support with pluggable SASL authentication mechanisms and optional TLS encryption
- Asynchronous model for session management and remote operation

Most of the features above are essential to the proper operation of a distributed application and yet non-trivial to implement. By offering them in the framework, we can significantly simplify the development of enterprise quality distributed applications.

DSP is officially registered with the [Internet Assigned Numbers Authority](#) under the service name dlpx-sp and port number 8415.

Currently, DSP supports Java language binding and provides a java based service framework for distributed applications.

Key Concepts

The foundation of DSP is built on top of a few key abstractions, namely, **exchange**, **task**, **nexus**, and **service**. For an overview of how DSP works and the features it provides, let's start with these abstractions.

An **exchange** refers to an application-defined protocol data unit which may be a request or a response. DSP supports the request-response pattern for communication. For each request sent, there is a corresponding response that describes the result of the execution. An application protocol is made up of a set of exchanges.

A **nexus** (a.k.a., session) refers to a logical conduit between the client and server application. In contrast, a transport connection (a.k.a., connection) refers to a “physical” link. A nexus has a separate naming scheme from the connection, which allows it to be uniquely and persistently identified independently of the physical infrastructure. A nexus has a different lifecycle than the connection. It is first established over a leading connection. After it comes into existence, new connections may be added and existing ones removed. It must have at least one connection to remain operational but may live on even after all connections are lost. Nexus lifecycle management actions, such as create, recover, and destroy, are always initiated by the client with the server remaining passive.

A nexus has dual channels, namely, the fore channel and the backchannel. The fore channel is used for requests initiated from the client to the server, and the backchannel from the server to the client. From a request execution perspective, the nexus is full-duplex and the channels are functionally identical, modulo the operational parameters that may be negotiated independently for each channel. A channel supports a number of features for request processing, such as ordered delivery, concurrent execution, remote cancellation, exactly-once semantics, and throughput throttling.

A **service** refers to a contract that consists of all exchanges (both the requests and the corresponding responses) defined in an application protocol. Given the full-duplex nature of request execution in DSP, part of the service is fulfilled by the server and the remaining by the client, where the client and server are from the nexus management perspective.

A **task** implements a workflow that typically involves multiple requests executed in either or both directions over the nexus. A task is a self-contained building block, available in the form of a sharable module including both the protocol exchanges and implementation, that can be easily integrated into other application protocols. A library of tasks may significantly simplify distributed application development by making it more of an assembly experience.

The following is a diagram that illustrates the key abstractions and how they are related to each other.

Security

As a network protocol, DSP is designed with security in mind from the onset. It supports strong authentication as well as data encryption. It follows a session-based authentication model which requires each connection to authenticate before it is allowed to join the session. Authentication is performed using the Simple Authentication and Security Layer (SASL) framework, a standard-based pluggable security framework. The currently supported SASL mechanisms include DIGEST-MD5, PLAIN with TLS, CRAM, and ANONYMOUS. Optionally, TLS encryption may be negotiated between the client and the server for data privacy.

Performance

DSP offers a number of features to enable the support for high-performance network applications. For example, it allows multiple requests to be exchanged in both directions simultaneously, which provides effective pipelining of data transfer to minimize the impact of network latency while ensuring the total ordering at the same time. It supports trunking that can effectively aggregate the throughput across multiple connections, which is crucial for a long fat network (LFN) and 10GigE. It also provides optional compression support which boosts performance over a bandwidth-limited network. We have observed, through both internal benchmarking and in customer environments, DSP-based applications delivering multi GigE in an ideal environment and getting a performance boost of as much as x10 in bandwidth-limited settings.

Resiliency

DSP automatically recovers from transient connection loss without any application involvement. It may also detect random data corruption on the wire and automatically recovers from it. In both cases, outstanding requests are retried once the fault condition is resolved.

DSP offers control over a remotely executing request. Once a request is initiated, the application may cancel it at any time before completion. In the rare event of a session loss, a new session creation request will be held until the old session has been reinstated. It ensures that we never leave any unknown or unwanted activities on the remote side and provides better predictability and consistency guarantees over an otherwise unreliable network.

Diagnosability

Application exceptions encountered during remote execution of a request are communicated back to the initiator through DSP. A standard Java API is used to facilitate the handling of remote exceptions that is in many ways identical to local ones.

DSP provides detailed information and statistics at the session level. The information may be used to examine the state of the session as well as diagnose performance problems. It is currently exposed via an internal support tool called the JMX tool.

Supported Applications

Replication is the first feature to take advantage of DSP. It has been rebuilt on top of DSP and shipping in the field since 3.1. In the latest release, a number of host-based applications, such as SnapSync, V2P, and Delphix connector, use DSP as well.

Related Topics


- [Network and DNS Management](#)
- [Optimal Network Architecture for the Delphix Engine](#)
- [CLI Cookbook: Network Performance](#)

Network Performance Test Tool Interface

Accessing the Network Performance Test Tool

To access the Network Performance Test Tool:

1. In the top navigation bar, click **Resources**.
2. Select **Network Performance Tool**.



The **Network Performance Tool** page will appear. There are two tabs: **Testing** and **History**.

Network Performance Tool

Testing History

Select Test Type

Latency - between this engine and a selected environment

Throughput - between this engine and a selected environment

DSP

Environment

Hawk ...

Address

bbdhep-AHCI-58503.dcenter.delphix.com

Number of requests to send

20

Request size (bytes)

16

[Run Test](#)

The Testing Tab

On the **Testing** tab, you can select one of three test types:

Test Type	Parameters Required
Latency	<ul style="list-style-type: none"> • Environment – If you selected an environment in the Datasets panel, this field will auto-populate. See below for instructions on changing the selected environment. • Number of requests to send • Request size (bytes)

Test Type	Parameters Required
Throughput	<ul style="list-style-type: none"> • Environment • Duration – in seconds • Number of connections • Direction – Select either Transmit or Receive • Block size (bytes)
Delphix Session Protocol	<ul style="list-style-type: none"> • Target Engine • Username • Password • Duration – in seconds • Block size (bytes) • Number of connections • Queue depth • Direction – Select either Transmit or Receive • Traffic Options

Selecting an Environment

For either a Latency or Throughput test, you must select an environment. If you selected an environment in the **Datasets** panel, this field will auto-populate. If you need to select a different environment:

1. Next to the **Environment** field, click the button with three dots.

Network Performance Tool

Testing History

Select Test Type

Latency - between this engine and a selected environment

Throughput - between this engine and a selected environment

DSP

Environment

Hawk

Address

bbdhcp-AHCI-58503.dcenter.delphix.com

The **Select Environment** screen will appear.

2. Select the environment you want to use in the test.

Environments	Hosts
source-l	source-l.dlpxdc.co
target-l	target-l.dlpxdc.co

3. Click **OK**.

Running and Canceling a Network Performance Test

To run a test:

1. Enter all the required parameters.
2. Click **Run Test**.
3. To cancel the test click **Cancel**.

The History Tab

In the **History** tab, you can view the results of all previous tests you have run.

1. Select the radio button for the type of test for which you want to see the results:
 - a. Latency, or
 - b. Throughput, or
 - c. Delphix Session Protocol
2. Optional: sort the tests by clicking one of the column headings.
3. Click the particular test for which you want to see the results.
The details of that test will appear.
4. Click **OK** to return to the **History** tab.

Related Topics

- [CLI Cookbook: Network Performance](#)
- [Delphix Session Protocol Test from Primary Engine to Replication Engine](#)

Delphix Session Protocol Test from Primary Engine to Replication Engine

Delphix uses the DSP protocol to communicate between primary and replication engines.

1. Login to Delphix Management application using an Engine administrator account.
2. Click the **Resource** menu and select **Network Performance Tool**.
3. Select the **Delphix Session Protocol** option.
 - a. Select Engine option to run the DSP test between a primary engine and a replication engine.
 - b. Provide hostname or IP address of the replication engine.
 - c. Provide admin user credentials for the replication engine. The user name must not include @DOMAIN or other qualifiers.
 - d. Enter a **Duration** in seconds (30 default).
 - e. Enter **Block Size** in bytes.
 - f. Enter the **Number of Connections** (optional).
 - g. Provide **Queue Depth**.
 - h. Select Direction - **Transmit** or **Receive**.
 - i. **Enter Traffic Options by selecting either Use compression or Use encryption.**
 - j. Click the **Run Test** button.

DELPHIX MANAGEMENT Manage Resources System Help

Network Performance Tool

Testing History

Select Test Type

Latency - between this engine and a selected environment

Throughput - between this engine and a selected environment

DSP

Engine - between this engine and a selected engine

Environment - between this engine and a selected environment

Environment

bbdhcp-1q6z-qar-52639-f39bd9ec.dcenter.delphix.com

Address

bbdhcp-1q6z-qar-52639-f39bd9ec.dcenter.delphix.com

Duration (seconds)

30

Block Size (bytes)

65536

Number of connections

0

Queue depth

32

Direction

Transmit

Receive

Traffic Options

Use compression

Use encryption

Run Test

4. View the test results.

The screenshot displays the 'Network Performance Tool' interface in the Delphix Management console. The top navigation bar includes 'DELPHIX MANAGEMENT', 'Manage', 'Resources', 'System', and 'Help'. The main content area is titled 'Network Performance Tool' and has two tabs: 'Testing' (active) and 'History'. A green checkmark icon indicates a successful 'DSP Test Result'. Below this, the target environment is listed as 'bbdhcp-1q6z-qar-52639-f39bd9ec.dcenter.delphix.com' and the target IP is '10.43.20.141'. Two large metrics are shown: '2.67 Gb/s THROUGHPUT' and '8 NUMBER OF CONNECTIONS'. A table of parameters and state is provided at the bottom.

PARAMETERS		STATE
Duration (seconds):	30 seconds	Completed
Direction:	Transmit	
Number of connections:	0	
Queue depth:	32	
Packet size:	65,536 bits	
Use compression:	false	
Use encryption:	false	
Start time:	Aug 11, 2020 2:20:01 PM	

Related Topics


- [Network Operations Using the Delphix Session Protocol](#)
- [CLI Cookbook: Network Performance](#)

Working with Dataset Performance

Accessing the Dataset Performance Graph

1. Log into the Delphix Management application.
2. In the **Resources** menu, select **Dataset Performance**.
3. Use the controls described below to view statistics and their related graphs.

General Graph Display and Controls



1 Timescale: Select from hour, 6 hours, day

2 View: View data by top 5 consumers or by the 75% percentile

3 Dateline: Displays timestamps of data points in the graph.

4 Timeline Selector: Specifies the start and end time for the currently displayed data. The range displayed is controlled by selecting the slider. Drag the slider to view statistics for the specific time.

5 Graph Legend: If more than one set of information is presented on the graph, Graph Legend displays a description and color for each set.

6

Dataset Table: Displays the information for each Dataset in a table format. Selecting a Dataset link takes you to the Dataset Status page.

7

Export: Exports the displayed table data.

Related Topics

- [Network and DNS Management](#)
- [Network Performance Configuration Options](#)
- [Network Performance Expectations and Troubleshooting](#)

Network Performance Expectations and Troubleshooting

Overview

Once you have run your [network performance tests](#) to each source and target environment, you should confirm that they meet expectations. Corporate networks commonly leverage 10Gbps "line speeds" between servers. Although networks are shared environments, switching infrastructure is typically used to isolate traffic between different hosts, with the goal of allowing each host to reach its potential. While many environments perform within the 70-90% range of line speed, in well-maintained environments we can see 90%+ line speed.

In some circumstances, you can exceed typical expectations by implementing more best practice recommendations – for example, putting two VMs on the same hypervisor (e.g. ESX) host, in the same chassis or blade enclosure, or when teamed (bonded) NIC cards are used for extra bandwidth. Furthermore, when Jumbo Frames are implemented correctly, they provide a substantial improvement.

If your network is not meeting expectations, you will almost always need help from someone whose role is focused on networking to help arrive at a root cause. You may also need to obtain temporary access to other VMs and/or physical hosts to isolate some issue(s). Additionally, you may need the help of Delphix Support or Professional Services to perform some tests from the Delphix Engine to systems that are not already connected environments within the Delphix product. (The CLI test will only work for environments connected to the Delphix Engine).

With assistance or not, you can do a number of things to narrow down the potential causes of poor performance between two systems by a process of elimination. Below are some high-level steps to consider. Keep in mind that network throughput will always represent the least performant component, so many of the steps below are intended to help isolate which component may be performing poorly.

Troubleshooting and information gathering questions in rough order of priority - record and share your answers

1. Have all source/target tuning settings been applied?
 - a. Is AIX in scope? LSO / LRO can have a significant impact, as can Jumbo Frames (see below)
2. What is the link speed on the hosts in question? Is NIC teaming / bonding / LACP in use?
 - a. Linux: `ethtool <device>`
 - b. Solaris: `dladm show-phys`
 - c. Windows: `wmic NIC where "NetEnabled='true'" get "Name,Speed"`
3. What are the test results with greater or fewer connections in parallel?
4. Can we test throughput to alternate servers? (See below)
5. What is the overall latency? What is the latency to each (OSI layer 3) hop? Is there one hop that consistency has a higher cost? (check the latency with ping and hops with traceroute)
6. How many devices (OSI layer 2) are in the path? Your network team will need to help you identify these devices.
 - a. Note: Only Layer 3 devices will show up when reviewing a traceroute however each layer two devices can impact traffic and each needs to be configured when implementing jumbo frames
 - b. Example devices in path to physical server: 1. virtual NIC -> 2. Virtual switch (ESX) -> 3. Chassis NIC -> 4. Rack switch -> 5. Core Switch -> 6. Rack Switch -> 7. Physical NIC
 - c. Example devices in path to virtual server: 1. virtual NIC -> 2. Virtual switch (ESX) -> 3. Chassis NIC -> 4. Rack switch -> 5. Chassis NIC -> 6. Virtual Switch (ESX) -> 7. Virtual NIC
7. What is the average network utilization on each hop? Is there congestion on any hop? (Network team will need to review)
8. Is QoS / VirtualConnect / 802.1p enabled? At what threshold will it engage? (Network team will need to review)
9. Is there a firewall or any deep packet inspection in the route? (Network team will need to review)
10. Are Jumbo Frames enabled on any or all hops? E.g. Delphix Engine, Virtual NIC, Virtual switch, and all hops down to the destination. (Network team will need to review)

- a. We have seen Delphix installations often benefit 10-20% from Jumbo frames, but certain platforms (such as AIX) can benefit much more dramatically
 - b. Note that JF enablement on two hosts without confirming all the network pieces are properly enabled will result in VERY poor performance
 - c. Test Jumbo Frames via with “Do Not Fragment” flag from the remote host to the Delphix Engine.
 - d. Note that typical MTU Jumbo Frame setting is 9000 bytes, although some vendors recommend a little above or below this.
 - e. The test below is at 8000, but you can test larger from there. Our goal is primarily to ensure that a number substantially larger than 1500 and somewhat close to the 9000 "de facto" standard is working.
 - f. Whenever two hosts connect, they perform a handshake called Path MTU negotiation, where they agree on the highest MTU they both support. This is how we avoid impact when communicating to hosts with differing MTUs.
 - g. Linux\$ ping -M do -s 8000 [Delphix_Engine_IP]
 - h. Windows> ping -f -l 8000 [Delphix_Engine_IP]
 - i. Solaris v10-# traceroute -F [Delphix_Engine_IP] 8000 ("Do Not Fragment" not supported by ping on Solaris until v11)
 - j. Solaris v11+# ping -s -D [Delphix_Engine_IP] 8972
11. Depending on the results above, a dedicated network or VLAN may help. Consider if that is an option for you. (Your network team will need to review)

Testing throughput testing to alternate servers

This will help isolate where a problem may be.

1. Delphix to Server A – already known
2. Delphix to Server B – physical; helps us see if there is a problem with the original server NIC or physical network settings
3. Server B to Server C – physical; helps us see if there is a problem with the Delphix server NIC or physical network settings
4. Delphix to Server D – virtual; helps us see if there is a problem with the virtual network or Delphix settings
5. Server D to Server E – virtual; same host; helps us see if there is a problem with the virtual network on the host
6. Server D to Server F – virtual; different host; helps us see if there is a problem with the virtual network

Conclusion

If you need further help, please contact Delphix Support or Professional Services to assist in getting the best performance possible from your environment.

Related Topics

- [Network and DNS Management](#)
- [Network Performance Configuration Options](#)
- [Working with Dataset Performance](#)

Determining the Delphix Server ID and Host Name

On occasion, it may be necessary to locate the **Delphix Server ID** and **Hostname**.

The Delphix Engine ID and Delphix Server ID are synonymous. The GUI currently uses "Server," and that is the terminology that will be used in this document.

The **Delphix Server ID** uniquely identifies each Delphix Engine. It is a 36-character hexadecimal string of the form **xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx**. You can view the Delphix Server ID in the Server Setup application, the Delphix Management application, or by using the Command Line Interface (CLI) method.

The **Hostname** is a name you assign. It typically matches the IP (DNS) name assigned to the Delphix Engine. The hostname can only be viewed by using the System Setup application or the CLI method.

Server Setup Application Method

Login to the Delphix Setup application with sysadmin-level credentials:

1. Access the Delphix Engine through the URL: `http://<Delphix Engine>/ServerSetup.html` where `<DelphixEngine>` is the DNS name or IP address of the Delphix Engine for which you wish to find the Delphix Server ID and hostname.
2. Enter a valid **Username**.
3. Enter a valid **Password**.
4. Click **Log In**.

On the **Dashboard** screen, there is a **System Summary** panel in the lower-left portion of the screen. The **Server ID** field displays the Delphix Server ID, in this example **564D39A8-5077-C9D0-9EFD-82E848EBDAB6**.

System Summary ⓘ	
Server ID	564D39A8-5077-C9D0-9EFD-82E848EBDAB6
Manufacturer	VMware
Model	Virtual_disk
Serial	6000c2907f0001 dd8000c240fee5b8c
Processor	2 x 2.90GHz
Memory	7.3GB
Features	XPP, MDD
Default locale	en-US

The Delphix Engine Hostname is located on the same screen, in the **Network** panel. The **Delphix Engine Hostname** field displays the hostname information.

Network ⓘ	Modify
Network Interface ens160	
Interface Configured Yes	
Jumbo Frames Enabled No	
IP Address Type DHCP	
IP Address 10.43.3.98	
Subnet Mask 255.255.0.0	
<hr/>	
Default Gateway 10.43.0.1	
DNS Domain Name delphix.com	
DNS Servers 172.16.101.11, 172.16.105.2	
Hostname js532.dcenter	

Delphix Admin Application

The Delphix Server Hostname is not available from this view but typically matches the IP (DNS) name assigned to the Delphix Engine.

Login to the Delphix Management application with delphix_admin level credentials:

1. Access the Delphix Engine through the URL: `http://<Delphix Engine>/Server.html`
where `<DelphixEngine>` is the DNS name of the IP address of the Delphix Engine for which you wish to find the Delphix Server ID.
2. Enter a valid **Username**.
3. Enter a valid **Password**.
4. Click **Log In**.

5. Under **System Summary**, the **Server ID** field displays the Delphix Server ID.

CLI Method

1. Use SSH to access your Delphix Engine: `ssh <userid>@<delphix_engine>`
where `<userid>` is a user ID with either `delphix_admin-` or `sysadmin-` level credentials.
2. Enter a valid password if prompted for one.
3. Enter `system ls`.
You will see an output similar to this example:

```
Properties

type: SystemInfo

apiVersion:

type: APIVersion

  major: 1

  micro: 0

  minor: 5

buildTimestamp: 2015-02-24T09:24:58.000Z

buildTitle: Delphix Engine 4.2.0.1

buildVersion:

  type: VersionInfo

  major: 4

  micro: 0

  minor: 2

  patch: 1

configured: true

currentLocale: en-US

enabledFeatures: XPP

hostname: delphix42.dcenter

installationTime: 2015-02-24T19:53:32.000Z

locales: en-US
```

```
memorySize: 3.99GB

platform: VMware with BIOS date 04/14/2014

processors:

  0:

    type: CPUInfo

    cores: 1

    speed: 2.40GHz

  1:

    type: CPUInfo

    cores: 1

    speed: 2.40GHz

productName: Delphix Engine

productType: standard

sshPublicKey: ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAszNCFnfzuiuK8dBdv6DNB+LrhVP1wRwc/
vXVrxrDlgyQTrqvEx4BKgHDZ2hnbAmqq2xXHR5Ah6WDSefo6u5B45JZc8qHpx8VZSza053IdMK9LEgoKPepmo7
JV3kVY9oHK9PngLm9tFK+hN7AUhcGTt68IHq54GWYQNBtx0kgSR5HtkkFhVfX2amFsHIsq1K96bgRkL0I5f3Sj
F4NnyElgBU9grGDajm9RXv+sz+F7h79AtFm0+W2Ymr5gQrdgh2vPyeFtG8G7rxnQx3qiRBY6lNqepBhitXnMY
SduGfW+fMjPv8T00J9ZLCfE7rjAgH7RxPybTfb4u70sm2krS8SgQ== root@delphix

storageTotal: 23.07GB

storageUsed: 2.00GB

uuid: 564d2f7c-b84f-8bd1-6f45-2060ac9b9a65
```

The **Delphix Server ID** is shown as the `uuid` property at the bottom of the output, and the **Hostname** is displayed in the `hostname` property.

4. Enter `exit` to leave the command-line interface.

Related Topics

- [Network and DNS Management](#)
- [Configuration](#)

Configuring Multiple DNS Domain Names in DNS Search List

This topic describes the steps to configure multiple DNS domain names in the DNS search list.

Procedure

Perform the following steps to configure multiple DNS domain names in the DNS search list.

1. Launch the Delphix Engine Setup interface using the sysadmin credentials.
2. Navigate to the **Network** widget and click **Modify**.
3. Under **DNS SERVICES**, use the **DNS Domain Name** and **DNS Servers** boxes for adding multiple DNS domain name configurations.

Network ✕

Configure network interfaces and services. Modifying the settings of network interfaces and default gateway may cause the Delphix Engine to be unreachable from the browser. It is recommended that such configuration be done from the Command Line Interface on the system console after a successful install.

NETWORK INTERFACES

ens192	Configured	Settings
---------------	------------	--------------------------

ROUTING

Default Gateway
10.43.0.1

DNS SERVICES

Set up DNS (Domain Name System) if you have a DNS Server in your environment. This will allow you to use names that will resolve into IP addresses when configuring components for your Delphix Engine.

Delphix Engine Hostname
sean6090.dcol2.delphix.com

DNS Domain Name ⓘ

DNS Servers ⓘ

Cancel
Save

Each domain name needs to be separated by a comma.

 **Note**

If .local DNS domains are in use, then you must add these explicitly to the list of DNS domains configured in order for name resolution to be successful. Multiple .local subdomains can be added as desired (for example, dev.company.local), or .local can simply be added to the DNS domain configuration to enable all .local domains to be successfully looked up in DNS. Multicast DNS is not currently supported by the Delphix Engine.

In order to understand if there is more than one domain name in the search list, check for "DNS Suffix Search List" from the output of `ipconfig /all` in the Windows Server:

```
C:\Program Files\Delphix\DelphixConnector\connector>ipconfig /all

Windows IP Configuration

Host Name . . . . . : 10-43-13-231
Primary Dns Suffix . . . . . : ad.delphix.com
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No
DNS Suffix Search List. . . . . : ad.delphix.com
                                delphix.com
                                dcenter.delphix.com
```

or check for "search" in `/etc/resolv.conf` on a Linux server:

```
[root@rhel62 ~]# cat /etc/resolv.conf
search delphix.com
nameserver 192.168.0.1
```

To Update DNS using the CLI:

1. Log into the CLI as sysadmin and navigate to `service > dns`.

```
ssh sysadmin@yourengine
> service
> dns
```

2. List the current DNS configuration and **update** to add new configurations.

```
> ls
> update
> set domain=xxx.xxx, xxx.xxx
```

3. **Commit** the action and verify the new list.

```
> commit
> ls
```

For example:

```
delphix> /service dns
delphix service dns> ls
```

```
Properties
type: DNSConfig
domain: delphix.com
node: (unset)
servers: 192.168.0.1

Operations
update

delphix service dns> update
delphix service dns update *> set
domain=delphix.com,one.delphix.com,two.delphix.com
delphix service dns update *> commit
delphix service dns>
delphix service dns> ls

Properties
type: DNSConfig
domain: delphix.com,one.delphix.com,two.delphix.com
node: (unset)
servers: 192.168.0.1

Operations
update

delphix service dns>
```

Related Topics

- [Network and DNS Management](#)

How to Change the IP Address of the Delphix Engine

Changing the IP Address

1. Stop all running VDBs by clicking the **Stop** button on the VDB card.
2. Disable all dSources.
3. You can use either the command-line interface or the Delphix Setup application to change the IP address of the Delphix Engine.
 - a. To use the command-line interface, follow the instructions described in [Setting Up Network Access to the Delphix Engine](#).
 - b. To use the Delphix Setup application, go to **System > Server Setup** in the Delphix Management interface, or click **Server Setup** in the Delphix Engine login screen.
 - i. In the **Network** panel, click **Modify**.
 - ii. Under **DNS Services**, enter the new IP address.
 - iii. Click **OK**.
4. Refresh all Environments by clicking the **Refresh** Symbol on the Environments screen.
5. Enable all dSources.
6. Start all VDBs by clicking the **Start** button on the VDB card.

Changing the IP Address via CLI

1. Stop all running VDBs by clicking the **Stop** button on the VDB card.
2. Disable all dSources.
3. Log into the Delphix CLI using your *sysadmin* account. You can find instructions on how to do this in the [Connecting to the CLI](#) article.

```
delphix> network
delphix network> setup
delphix network interface> list
NAME
vmxnet3s0
delphix network interface> select vmxnet3s0
delphix network interface 'vmxnet3s0'> get
  type: NetworkInterface
  name: vmxnet3s0
  addresses:
    0:
      type: InterfaceAddress
      address: 10.1.2.3/24
      addressType: STATIC
      enableSSH: true
      state: OK
  dataNode: DATA_NODE-34
  device: vmxnet3s0
  macAddress: 0:c:29:32:96:a3
  mtu: 1500
  mtuRange: 60-9000
  reference: NETWORK_INTERFACE-vmxnet3s0-DATA_NODE-34
  state: OK
```

4. Run the update command and update the address to the new IP address for the Delphix Engine.

```
delphix network interface 'vmxnet3s0'> update
delphix network interface 'vmxnet3s0' update *> edit addresses.0
delphix network interface 'vmxnet3s0' update addresses.0 *> get
Properties
  type: InterfaceAddress
  address: 172.16.151.154/24
  addressType: STATIC
  enableSSH: true

delphix network interface 'vmxnet3s0' update addresses.0 *> set address=10.1.2.
4/24
delphix network interface 'vmxnet3s0' update addresses.0 *> get
  type: InterfaceAddress (*)
  address: 10.1.2.4/24 (*)
  addressType: STATIC (*)
  enableSSH: true (*)
```

5. Commit the operation.

```
delphix network interface 'vmxnet3s0' update addresses.0 *> commit
delphix network interface 'vmxnet3s0'> get
  type: NetworkInterface
  name: vmxnet3s0
  addresses:
    0:
      type: InterfaceAddress
      address: 10.1.2.4/24
      addressType: STATIC
      enableSSH: true
      state: OK
  dataNode: DATA_NODE-34
  device: vmxnet3s0
  macAddress: 0:c:29:32:96:a3
  mtu: 1500
  mtuRange: 60-9000
  reference: NETWORK_INTERFACE-vmxnet3s0-DATA_NODE-34
  state: OK
```

6. Re-enable the VDBs and dSources running from the engine.

Related Topics

- [Network and DNS Management](#)

How to Change the Hostname of the Delphix Engine

Changing the Hostname

Perform the following steps to change the hostname of a Delphix engine.

Note

Currently, it is only possible to change the hostname via the Command Line Interface (CLI).

1. Stop all running VDBs by clicking the **Stop** button on the VDB card.
2. Disable all dSources.
3. Log into the Delphix CLI using your **sysadmin** username and password. If you are using the same interface that you connected to the CLI on, it will interrupt the connection. Therefore, it is recommended to log in to the CLI on the console.

```
ssh sysadmin@yourdelphixengine
```

4. Run the following commands.

```
dlpx01> network
dlpx01 network> setup
dlpx01 network setup> ls
Properties
  defaultRoute: 192.168.0.1
  dhcp: false
  dnsDomain: plb.internal
  dnsServers: 192.168.0.111,198.142.152.164,198.142.152.165
  hostname: dlpx01
  primaryAddress: 192.168.0.109/24

Operations
update
dlpx01 network setup> update
dlpx01 network setup update *> set hostname=newhostname
dlpx01 network setup update *> commit
```

5. If DHCP is being used, the Delphix Engine will expect the Hostname to be provided by the DHCP server. As such, there will be no property 'hostname' to update. This process requires changing the IP addressing configuration from DHCP to static address configuration. To proceed you will need to take the following action (making sure that there is no conflict between your DHCP server and the changes you are implementing).

```
dlpx01 network setup> update
dlpx01 network setup update *> set dhcp=false
dlpx01 network setup update *> set hostname=newhostname
dlpx01 network setup update *> commit
newhostname network setup> update
newhostname network setup update *> set dhcp=true
```

```
newhostname network setup update *> commit
newhostname network setup> ls
Properties
  defaultRoute: 192.168.0.1
  dhcp: true
  dnsDomain: plb.internal
  dnsServers: 192.168.0.111,198.142.152.164,198.142.152.165
  hostname: newhostname
  primaryAddress: 192.168.0.109/24
```

6. Re-enable the VDBs and dSources running from the engine.

Related Topics

- [Network and DNS Management](#)

How to Change the DNS Server of the Delphix Engine

When DHCP is enabled on the Delphix Engine, the DNS server cannot be changed in the GUI. In this case, the DNS server can be changed in the command-line interface (CLI). If the IP address, net mask, or any other properties of the network interface itself are not changed, then VDBs do not need to shutdown. Login via Hypervisor Console for the Delphix Engine virtual machine is optional.

Use the commands below and enter the address of a server or a comma-separated list of addresses in replace of the example servers.

```
ssh -l sysadmin DelphixEngineName
Password:
DelphixEngineName> service dns
DelphixEngineName service dns> update
DelphixEngineName service dns update *> set servers=1.1.1.1,2.2.2.2
DelphixEngineName service dns update *> commit
DelphixEngineName service dns> exit
```

If DHCP is disabled and the address is static, the DNS server can be updated via the GUI.

1. Login to **Delphix Setup** with the appropriate credentials.
2. Locate the **Network** tile in the Dashboard and select **Modify**.

The screenshot shows the Delphix Setup web interface. The top navigation bar includes 'DELPHIX SETUP', 'Dashboard', 'Preferences', 'Support Bundle', and 'Help'. A user is logged in as 'sysadmin'. The main content area is divided into several sections:

- System Summary:** Displays hardware and system information such as Server ID, Manufacturer, Model, Serial, Processor, Memory, Features, and Default locale.
- Network:** The active configuration window for the 'en' network interface. It shows:
 - Network Interface: en
 - Interface Configured: Yes
 - MTU: 1500
 - IP Address Type: STATIC
 - IP Address: 10
 - Subnet Mask: 25
 - Default Gateway: 10
 - DNS Domain Name: delphix.com
 - DNS Servers: 17
 - Hostname: ryan
- User-click Analytics:** Disabled.
- SMTP Server:** Disabled. A warning message states: 'SMTP not configured - functionality reduced: Not able to send or receive events and notifications.'
- Registration:** View button.
- Engine UUID:** (Field is empty)
- Status:** (Field is empty)

3. Scroll down in the **Network** window.

4. Use the **DNS Domain Name** and **DNS Servers** boxes for new configurations.

Network ✕

ROUTING

Default Gateway
10

DNS SERVICES

Set up DNS (Domain Name System) if you have a DNS Server in your environment. This will allow you to use names that will resolve into IP addresses when configuring components for your Delphix Engine.

Delphix Engine Hostname
ryan

DNS Domain Name ⓘ

DNS Servers ⓘ

Cancel Save

Configuring a Second Network Interface

This topic describes how to configure a static IP address on a second network interface.

Procedure

Perform the following steps to configure a second network interface.

1. Launch the Delphix Engine Setup interface using the sysadmin credentials.
2. Navigate to the **Network** widget and click **Modify** to view the available network interfaces, and select the new interface to be configured.
3. Click the **Settings** button next to the network interface that you want to configure.


Network ✕

Configure network interfaces and services. Modifying the settings of network interfaces and default gateway may cause the Delphix Engine to be unreachable from the browser. It is recommended that such configuration be done from the Command Line Interface on the system console after a successful install.

NETWORK INTERFACES

ens256	Not configured	Settings
ens224	Not configured	Settings
ens192	Configured	Settings

4. The **Network Interface Settings** screen appears. Select the checkbox before **Enabled** to enable the network.



The screenshot shows a dialog box titled "Network Interface Settings" with a close button (X) in the top right corner. The interface name "ens224" is displayed. Below it, there are two options: "Enabled" (unchecked) and "DHCP" (unchecked). The "Static" option is selected with a radio button. A group box contains two input fields: "IP Address" and "Subnet Mask", both containing placeholder text. Below this group box is an "MTU" field with a value of "1500" and an information icon (i). At the bottom right, there are two buttons: "Cancel" and "Save".

5. Select one of the following: **DHCP** or **Static** and enter the **IP address** and **Subnet Mask** address in the respective fields.
6. **MTU**: Enter a value for the MTU field. This is the maximum size in bytes of a packet that can be transmitted on this interface.
7. Click **Save** to save the settings.

NFSv4 Configuration

Overview

This article shows which OS versions support NFSv4. NFSv4 is enabled by default, but the target/staging hosts require additional configuration changes before it can be used.

Redhat/SLES: "NFSv4 Only - Enabling Recover Lost Locks" section under [Linux/Redhat/CentOs](#).

AIX: [IBM AIX](#) article.

OS version	NFSv4 support
RHEL 6.3 or later	Supported
SLES 11.4 or later	Supported
AIX 7.1, 7.2	Supported

If these target host configurations are not set, the engine will choose NFSv3. The provided reason (e.g., "OS not supported") that is shown for choosing NFSv3 over NFSv4 can be observed under the Status tab of the selected VDB in the Delphix Engine GUI. Additional information is provided in this tab for datasets and current values.

Dataset Status Tab NFS Reasons

NFS Reason	Notes
Default	NFSv4 is used by default.
Old RedHat	NFSv3 is used for versions of RedHat older than 6.3.
Unsupported OS	NFSv3 is used for HP-UX and Solaris.
DNFS	dNFS is in use and does not support NFSv4.
Tunable Override	The nfs.version tunable is set to force NFSv3.
Configuration Override	Appears with default operation, implies that a configuration parameter has changed.
No Recover Lost Locks	The required Linux NFS client 'recover_lost_locks' setting is not enabled.
Unsupported Data Platform	NFSv3 is used with AppData.

NFS Reason	Notes
Incomplete v4 Config	<code>nfsrgyd</code> daemon is not running or reverse DNS lookup is missing.

Related Topics

- [Target Host OS and Database Configuration Options](#)

Capacity and Resource Management

Delphix will be responsible for managing many different data sources and data types. As such, it is critical to understand how to manage your capacity and storage resources within each Delphix engine. Learn storage and quota best practices, as well as the different options to optimally manage capacity.

This section covers the following topics:

- [An Overview of Capacity and Performance Information](#)
- [Setting Quotas](#)
- [Deleting Objects to Increase Capacity](#)
- [Adding and Expanding Storage Devices](#)
- [Delphix Storage Migration](#)
- [Managing Source Data](#)
- [An Overview of Held Space](#)

An Overview of Capacity and Performance Information

This topic describes the Delphix Engine performance reservoir and capacity threshold warnings and various ways to obtain information about capacity and resource usage for the Delphix Server.

The Performance Reservoir and Capacity Threshold Warnings

In order to obtain the best performance and continued operations, the Delphix Engine requires a certain amount (minimum 512GB) of free space of the total quota for storage space. As storage capacity approaches this threshold, the following system faults occur:

- When **85%** of the total storage quota is reached or **1536GB** of free space is remaining (whichever is less), a **Warning** fault is triggered. You can resolve this fault by deleting objects in the Delphix Engine, adding storage, or changing policies, as described in the topics [Adding and Expanding Storage Devices](#), [Deleting Objects to Increase Capacity](#). Additionally, refreshing target databases will clear the space the engine uses to track changes over time for each DB.

 When a **Warning** fault is first raised, it will not impact the functioning of the system. However, there may be an impact to Oracle LogSync and SQL Server Validated sync when a **Warning** fault is seen as storage usage decreases following a **Critical** storage usage fault.

- When **90%** of the total storage quota is reached or **1024GB** of free space is remaining (whichever is less), a **Critical** fault is triggered. You cannot ignore or resolve this fault. This fault will have a significant impact on the system's behavior such as:
 - All pending link, sync, refresh, and provisioning processes will be canceled, and no new operations can be initiated
 - Scheduled replication processes will first check for capacity on the target engine and hold data (replication currently in progress to the engine will not be halted)
 - Policy operations such as SnapSync, snapshot, and refresh are suspended for **all** databases
 - dSources stop pulling in new changes. LogSync is suspended for all **Oracle dSources**. Validated sync is disabled for **SQL Server** dSources.
 - No virtual database (VDB) snapshots can be taken.
- When **95%** of the total storage quota is reached or **512GB** of free space is remaining (whichever is less), a second **Critical** fault is triggered. This fault will have a significant impact on the system's behavior and certain dSources and VDBs will stop in order to maintain data integrity. For example, SQL Server VDBs will shut down.

As the free space of the system improves, the following functionalities can be automatically or manually resumed.

- When the system falls below **95%** of the total storage quota, you can manually start SQL Server VDBs that had stopped
- When the system falls below **90%** of the total storage quota:
 - SQL Server VDBs that had stopped will automatically start
 - New link, sync, refresh and provisioning operations are allowed
 - Policy operations such as SnapSync, Snapshot, and Refresh resume for all databases
- When the system falls below **85%** of the total storage quota
 - dSources start pulling in new changes from their corresponding data sources. LogSync is resumed for **Oracle dSources**. Validated sync is enabled for **SQL Server** dSources.

For more information, see [Setting Quotas](#).

Ways to View Capacity Usage

You can access capacity and performance information for the Delphix Engine through several different means, including the **Dashboard** view, and the **Storage Capacity** screen.

The Dashboard View

You can access the **Dashboard** view in the **Delphix Management** application by clicking **Dashboard** in the Manage menu. Note that the Dashboard view provides only summary information about capacity and performance. You must access the **Storage Capacity** and **Dataset Performance** screens in the **Resources** menu to manage storage space and database objects.

The Dashboard view provides more detailed information about the overall performance of the Delphix Engine:

- **Storage Capacity** - the amount of physical storage available and what is currently used
- **TimeFlow Ratio** - see above
- **VDB Ratio** - a measure of the amount of physical space that would be occupied by the database content against the amount of storage occupied by that same database content as VDBs.
- **Dataset Performance** - the amount of network bandwidth available and the amount that VDBs are currently utilizing, as well as information about specific VDB network usage

The Storage Capacity Screen


You can access the **Storage Capacity** screen through the **Resources** menu in the **Delphix Management** application. The **Storage Capacity** screen now has an improved format the interface makes it easier to identify which datasets are consuming the most space on your Delphix Engine.

Major Additions in 6.0.6:

- Users can now see what's pinning a held space and batch delete the dependencies to free it up.

What's Changed 6.0.6:

- Snapshot Delete dialog now shows a tree of "prerequisites" and "dependencies" that must have operations performed on them / be deleted in order to enable the deletion of the locked root snapshot.
- The Snapshot Delete dialog now automates the steps.
- Held Space shown in the table now represents deadbeat Timeflows rather than Containers



For more information please refer to [Using and Understanding The Storage Capacity Screen](#).

Related Topics

- [Using and Understanding The Storage Capacity Screen](#)
- [Deleting Objects to Increase Capacity](#)
- [Setting Quotas](#)
- [Adding and Expanding Storage Devices](#)

Using and Understanding The Storage Capacity Screen

The **Storage Capacity** screen can be reached via the **Capacity** item in the **Resources** menu. This screen shows how storage capacity is allocated for dSources, VDBs, and Snapshots, and permits storage space to be reclaimed through the deletion of objects and Snapshots.

Storage Capacity 1 Reload Data

2 **18.76GB Remaining of 21.80GB Total**

● Reserved 3.27GB ● Held 74.49MB ● Your data 3.04GB

3 Datasets Snapshots Held Space Received Replicas How it Works

4 View: All Status: All Type: All Search 5

6 Dataset	Size	Group	Delete	Action
> dbdhcp2:gli Oracle 11.1.0.7.0	482.95MB	gli	🗑️	⋮
> dbdhcp1 Oracle 10.2.0.5.0	217.53MB	Untitled	🗑️	⋮
> C4 Oracle 10.2.0.5.0	169.80MB	Untitled	🗑️	⋮
> C3 Oracle 10.2.0.5.0	166.98MB	Untitled	🗑️	⋮
> C1 Oracle 10.2.0.5.0	166.14MB	Untitled	🗑️	⋮
> Child VDB Oracle 11.1.0.7.0	135.39MB	gli	🗑️	⋮
> VDB-1 Oracle 11.1.0.7.0	134.11MB	gli	🗑️	⋮
> Parent VDB Oracle 11.1.0.7.0	133.97MB	gli	🗑️	⋮
> C5 Oracle 10.2.0.5.0	92.25MB	Untitled	🗑️	⋮
> C6 Oracle 10.2.0.5.0	91.85MB	Untitled	🗑️	⋮

Items per page: 15 1 - 11 of 11 |< < > >|

Storage Capacity Screen

1

Select the Reload Data button to clear your capacity cache and refresh the page. If you are clearing the cache of a large Delphix Engine, this operation could take several minutes. During this operation, the capacity page is not available.

2

The usage Graph at the top of the page provides a visual representation of the distribution of storage consumption across your engine.

- **Reserved** - The amount of reserved space is hard-coded to be 10% (maximum up to 1TB) of the system capacity and cannot be changed. The system has a 90% consumption cap.
- **Held** - Total held space size

- **Your data** - User data (snapshot, log, etc.); Includes shared space

 Currently, this only includes dataset space. It doesn't include space of the MDS.

3

Click a tab to view one of the following tables:

- **Datasets** - this tab shows a list of datasets on the engine ordered by the size of the dataset. You can filter the list by whether the dataset is VDB or dSource, whether the dataset is locked or unlocked, and the type of dataset. The drop-down drawer of each dataset shows some metadata about the dataset. If the dataset has descendants, the drop-down drawer will also show the list of **first-level** descendants. You can use the Action menu (...) to delete or refresh a dataset. If a dataset has descendants, the delete column will show a lock icon which opens a dialog explaining why this dataset is locked. Replicated datasets are excluded.
- **Snapshots** - rows are always sorted by descending size with a link to the dataset page for the Snapshot's parent dataset. Snapshots created from Replicas are excluded. The drop-down drawer for each snapshot shows the latest snapshot date of the dataset that this snapshot is from. If a snapshot has any descendants or can't be deleted due to descendants, the delete column will show a locked icon indicating this snapshot is locked. Clicking on the lock icon will display the snapshot why-locked dialog. Freshly created Snapshots are not displayed as this table is not updated automatically.
- **Held Space** - provides a list of Held Space on the engine. You can search the held spaces by their names (in this case their references). As shown below users can now see what's pinning a held space and batch delete the dependencies to free it up.

Delete Held Space "STORAGE_CONTAINER-208" ✕

This object is locked due to dependencies. To delete the snapshot, the actions listed under each "Prerequisites" section must be performed and the items listed under each "Dependencies" section must be deleted. You may deselect the items that you do not wish to delete (deselecting an item will cause its ancestors to be deselected). Once selection is finalized, clicking "Delete" will delete the selected items. Note that the total size displayed at the bottom is an estimate. The actual size recovered may be less, since the space on disk may be shared by other snapshots.

Selection	Name	Size
<input checked="" type="checkbox"/>	Held Space (STORAGE_CONTAINER-208)	650.38MB
	Dependencies	--
<input checked="" type="checkbox"/>	Dataset (VDBO_GN5)	441.59MB

2 Items Selected / ~1.07GB

Copy to Clipboard
Cancel
Delete

- **Received Replicas** - provides a list of Namespaces and by default is sorted by descending size. This tab is updated automatically (except the size). The drop-down drawer lists groups inside of each replicas sorted by size.
- **How it Works** - provides general information about Storage, Retention Policy and Manual deletion.

4

Filters available are table-specific and do not apply to the content inside a drawer

- **View** - Dataset type (VDB, dSource, vFiles, etc.)
- **Status** - locked/unlocked for deletion
- **Type** - database type (Oracle, SQL Server, AppData, etc.)
- **Search** - search by name (case insensitive)

5

Select  to export the information provided in the grid to a.csv file. Note: Only the page you are viewing will be exported.

6

The **table** section displays a set of information about objects tracked by the Delphix Engine. The information displayed varies depending on the selected tab.

Datasets	
Dataset	Dataset’s type as an icon and name. Datasets created from replicas are excluded and deleted Datasets have a strike-through their name.
Size	Space actually used by the object.
Group	The name of the Dataset Group.
Delete	Either a Trashcan or a locked icon will be displayed. The Trashcan opens a confirm delete dialog, for Oracle users can select to provide credentials. A Lock icon as shown below opens a dialog explaining why the selected object cannot be deleted.
Action	Refresh actions that can be taken against the Dataset. <ul style="list-style-type: none"> • Refresh N/A • Refresh Dataset
Snapshots	
Timestamp	The timestamp reflecting when the snapshot was taken.
Size	Space actually used by the snapshot.
Dataset	The name of the Dataset.
Group	The name of the Dataset Group.

Delete	Either a Trashcan or a locked icon will be displayed. The Trashcan opens a confirm delete dialog. A Lock icon as shown below opens a dialog explaining why the selected object cannot be deleted.
Held Space	
Object	The name of the object.
Size	Space actually used by the object.
Delete	For Held Space every row has a lock icon for opening up the Why-locked dialog.
Received Replicas	
Namespace	Once replication is complete, the target Engine will create a received replica, also known as a namespace.
Size	The amount of space used by the Namespace.

7


Dataset Table or Snapshot drawer provides information on dependent or descendant. The inner dependency table is direct children only and has the same behavior as the outer table. Filter or sorting from the outer table does not apply to the inner table.

The screenshot displays the Delphix Engine interface for a dataset named 'dbdhcp2:gli' (Oracle 11.1.0.7.0) with a size of 482.95MB. The 'Dependent Datasets' table is shown below, listing two datasets: 'VDB-1' (134.11MB) and 'Parent VDB' (133.97MB), both in the 'gli' group. Below the table, summary statistics for the current copy are provided:

Current Copy Size	443.55MB
DB Log Size	19.75MB
Total Snapshot Size	19.57MB
Temp File Size	0.00B
Shared Snapshot Space	0.00B

The interface also shows a bottom section for 'dbdhcp1' (Oracle 10.2.0.5.0) with a size of 217.53MB and an 'Untitled' group.

The **Delete Dataset dialog** - this dialog provides information on why a dataset cannot be deleted. It displays **all** descendants of a dataset, as well as any self-service objects on any of the descendant datasets. You can trace the chain of descendants from the current dataset all the way to leaf datasets. Click **Copy to Clipboard** to copy the list of descendants to clipboard.



The **Delete Snapshot Dialog** - The new snapshot dialog provides the complete steps on how to unlock the locked snapshot for deletion. It also allows users to select descendants and batch delete them. The dialog shows a tree of **prerequisites** and **dependencies** that must have an operation performed or be deleted in order to enable the deletion of the locked root snapshot. **Prerequisites** refer to operations that must be performed before the object it's associated with can be deleted. Examples of prerequisite operations include: removing a keep-forever retention policy from a snapshot, deleting a Self-Service bookmark, or refreshing a dataset. **Dependencies** refer to objects that must be deleted before the parent object can be deleted. There is also a special **All (other) Snapshots** dependency item. This item represents all children snapshots of the parent Timeflow, which will also need to be removed, minus those with a dedicated row in the table.

Delete Snapshot "Jan 7, 2021 2:05:41 AM -05:00" ✕

Selection	Name	Size
<input checked="" type="checkbox"/>	Snapshot (Jan 6, 2021 11:05:41 PM)	13.00KB
<input type="checkbox"/>	Prerequisites	-
<input checked="" type="checkbox"/>	Take a snapshot on container (dbdhcp1)	-
<input type="checkbox"/>	Dependencies	-
<input checked="" type="checkbox"/>	Timeflow (DB_PROVISION@2021-01-06T23:08:58) on dataset (Vdbd_JXM)	0.00B
<input type="checkbox"/>	Prerequisites	-
<input checked="" type="checkbox"/>	Refresh container (Vdbd_JXM)	-
<input type="checkbox"/>	Dependencies	-
<input checked="" type="checkbox"/>	All (other) Snapshots	573.50KB

0 Items Selected / ~0.00B

Related Topics

- [Capacity and Resource Management](#)

Understanding Delphix Disk Usage

Within the Delphix Timeflow, incremental restore points are called Snapshots. Delphix Snapshots use a shared block architecture to provide all the functionality of daily full backups using only a tiny percentage of the storage for full backups. Because shared block architectures are not common, a "common sense" understanding of how storage is associated with a Snapshot can be misleading.

- The size of a Snapshot is defined to be the size of changes that are unique to that Snapshot. In other words, changed blocks that are only associated with that one Snapshot.
- The latest Snapshot will always have 0 sizes initially, as there are no changes associated with it, nothing has changed since the Snapshot was taken, so there is no space used by unique changes.
- A block can be shared by multiple Snapshots if the block has not changed between the creation of those Snapshots. Any blocks that are shared amongst multiple Snapshots are accounted for in the shared Snapshot Space total.

The following screenshot provides an example of disk usage.

Dataset	Size	Group	Delete	Action										
✓ CDOMSHSRD97A <small>Oracle 18.3.0.0.0</small>	1.14GB	test	🗑️	⋮										
<div style="border: 1px solid #ccc; padding: 5px;"> <table> <tr> <td>Current Copy Size</td> <td>1.01GB</td> </tr> <tr> <td>DB Log Size</td> <td>18.80MB</td> </tr> <tr> <td>Total Snapshot Size</td> <td>118.12MB</td> </tr> <tr> <td>Temp File Size</td> <td>0.00B</td> </tr> <tr> <td>Shared Snapshot Space</td> <td>Not available</td> </tr> </table> </div>					Current Copy Size	1.01GB	DB Log Size	18.80MB	Total Snapshot Size	118.12MB	Temp File Size	0.00B	Shared Snapshot Space	Not available
Current Copy Size	1.01GB													
DB Log Size	18.80MB													
Total Snapshot Size	118.12MB													
Temp File Size	0.00B													
Shared Snapshot Space	Not available													
> CDOMSHTG661F <small>Oracle 18.3.0.0.0</small>	1.02GB	test	🗑️	⋮										
> VDBO_OBA <small>Oracle 18.3.0.0.0</small>	1,009.71MB	test	🔒	⋮										
> DBOMSRBBD0C8C <small>Oracle 18.3.0.0.0</small>	922.83MB	test	🔒	⋮										
> CDOMLOSRB5A3 <small>Oracle 18.3.0.0.0</small>	891.08MB	Untitled	🗑️	⋮										
> CDOMLOTG4F5E <small>Oracle 18.3.0.0.0</small>	859.85MB	Untitled	🗑️	⋮										

- **Current Copy Size:** The current copy size is the amount of space used on the Delphix Engine by the VDBs data but across all the Time flows.
- **DB Log Size:** The DB log size is the size for archive log files.
- **Temp File Size:** This field is currently not active (for future use).
- **Container Size:** The container size is the size of the whole dataset.
- **Total Snapshot Size:** The total snapshot size is the size of all snapshots combined, including anything shared.
- **Shared snapshot space:** The shared snapshot space is the space shared by snapshots.

Related Topics

- [Capacity and Resource Management](#)
- [An Overview of Capacity and Performance Information](#)

Reviewing Historical Capacity From the CLI

You can retrieve historical capacity data through the Delphix Command Line Interface (CLI).

If your Delphix Engine is a new engine you would log in as Engine Admin and ssh admin@your engine. Engines created before 5.3.1 and upgraded to 5.3.1 or later will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then disabling delphix_admin.

Process

1. Log in as Delphix Admin (or an account with admin privileges)

```
ssh delphix_admin@yourengine
```

2. Navigate to capacity system historical

```
delphix > capacity system historical
```

3. Then you can list a start and end date for the utilization

```
delphix capacity system historical > list startDate=<time> endDate=<time>
```

For example, looking at the system space utilization, we can show 8 hours starting Sep 1 with the following:

```
delphix> capacity
delphix capacity> system
delphix capacity system> historical
delphix capacity system historical> list startDate=2016-09-01T00:00:00.000Z
endDate=2016-09-01T08:00:00.000Z
TIMESTAMP SOURCE.ACTUALSPACE VIRTUAL.ACTUALSPACE
2016-09-01T00:25:11.359Z 5.05TB 9.58TB
2016-09-01T00:55:28.869Z 5.06TB 9.44TB
2016-09-01T01:25:42.940Z 5.06TB 9.46TB
2016-09-01T01:56:14.585Z 5.06TB 9.38TB
2016-09-01T02:26:50.893Z 5.06TB 9.38TB
2016-09-01T02:57:15.987Z 5.06TB 9.32TB
2016-09-01T03:27:35.381Z 5.06TB 9.34TB
2016-09-01T03:57:54.657Z 5.06TB 9.35TB
2016-09-01T04:28:12.099Z 5.06TB 9.35TB
2016-09-01T04:58:22.028Z 5.06TB 9.35TB
2016-09-01T05:28:33.680Z 5.06TB 9.34TB
2016-09-01T05:58:46.666Z 5.06TB 9.27TB
2016-09-01T06:28:57.181Z 5.06TB 9.21TB
2016-09-01T06:59:39.567Z 5.06TB 9.21TB
2016-09-01T07:30:03.527Z 5.06TB 9.20TB
2016-09-01T07:50:18.548Z 5.06TB 9.20TB
```

We can also set the output detail level to different granularity, which is based on **seconds**.

The following example will show the data for **each day** (86400 seconds) from Sep 1 to Oct 1:

```
delphix capacity system historical> list startDate=2016-09-01T00:00:00.000Z
endDate=2016-10-01T08:00:00.000Z resolution=86400
TIMESTAMP SOURCE.ACTUALSPACE VIRTUAL.ACTUALSPACE
2016-09-01T00:25:11.359Z 5.05TB 9.58TB
2016-09-02T00:34:16.844Z 5.03TB 8.45TB
2016-09-03T00:37:50.372Z 5.10TB 8.30TB
2016-09-04T00:32:22.877Z 4.68TB 8.79TB
2016-09-05T00:34:37.715Z 4.64TB 8.56TB
2016-09-06T00:37:27.480Z 4.65TB 8.44TB
2016-09-07T00:31:17.808Z 4.70TB 8.60TB
2016-09-08T00:33:34.219Z 4.74TB 8.70TB
2016-09-09T00:43:40.760Z 4.81TB 8.57TB
2016-09-10T00:48:27.222Z 5.14TB 8.95TB
2016-09-11T00:50:41.843Z 4.90TB 8.72TB
2016-09-12T00:50:33.215Z 4.92TB 8.82TB
2016-09-13T00:48:06.350Z 4.93TB 8.72TB
2016-09-14T00:42:36.904Z 5.34TB 8.79TB
2016-09-15T00:48:58.580Z 0B 0B
2016-09-16T00:43:12.565Z 5.64TB 9.23TB
2016-09-16T21:46:06.333Z 5.74TB 9.29TB
2016-09-18T19:41:29.692Z 5.74TB 9.29TB
2016-09-19T19:38:35.268Z 5.47TB 7.94TB
2016-09-20T19:39:07.809Z 5.57TB 8.50TB
2016-09-21T19:42:32.602Z 5.64TB 9.04TB
2016-09-22T19:37:53.507Z 5.82TB 8.77TB
2016-09-23T19:37:43.373Z 5.88TB 8.96TB
2016-09-24T19:29:40.625Z 5.78TB 9.03TB
2016-09-25T19:32:30.592Z 5.80TB 8.91TB
2016-09-26T19:26:43.971Z 5.81TB 8.86TB
2016-09-27T19:33:26.939Z 5.87TB 9.44TB
2016-09-28T19:33:01.483Z 5.94TB 9.44TB
2016-09-29T19:43:55.451Z 6.15TB 9.08TB
2016-09-30T19:37:38.462Z 6.19TB 9.72TB
2016-10-01T07:35:53.644Z 5.87TB 8.98TB
```

The following query will show the data for each week (604800 seconds), for a few months:

```
delphix capacity system historical> list startDate=2016-06-01T00:00:00.000Z
endDate=2016-10-01T08:00:00.000Z resolution=604800
TIMESTAMP SOURCE.ACTUALSPACE VIRTUAL.ACTUALSPACE
2016-06-01T00:16:19.391Z 3.36TB 5.96TB
2016-06-08T00:16:04.620Z 2.79TB 6.11TB
2016-06-15T00:16:35.471Z 2.95TB 4.84TB
2016-06-22T00:22:57.619Z 0B 0B
2016-06-29T00:28:48.350Z 3.38TB 6.53TB
2016-07-06T00:38:20.643Z 3.38TB 7.33TB
2016-07-13T00:36:38.876Z 3.30TB 8.05TB
2016-07-20T00:50:24.318Z 3.33TB 8.95TB
2016-07-27T00:49:06.488Z 3.36TB 9.28TB
2016-08-03T01:03:59.577Z 3.59TB 9.25TB
2016-08-10T01:12:21.949Z 3.81TB 8.80TB
2016-08-17T01:20:44.878Z 4.08TB 9.89TB
2016-08-24T01:21:01.080Z 4.23TB 9.00TB
2016-08-31T01:29:51.907Z 4.45TB 9.48TB
2016-09-07T01:31:56.383Z 4.70TB 8.67TB
2016-09-14T01:43:21.456Z 5.34TB 8.87TB
2016-09-21T01:42:34.318Z 5.58TB 8.84TB
2016-09-28T01:28:10.313Z 5.88TB 9.47TB
2016-10-01T07:35:53.644Z 5.87TB 8.98TB
```

Related Topics

- [Capacity and Resource Management](#)


Setting Quotas

This topic describes how to set quotas for database objects.

1. Log into the **Delphix Management** application.
2. Select **Manage > Datasets**.
3. Select the **Configuration** tab for the object you want to set a quota for.
4. Select the **Policies** tab and from the Quota tile select **Customize this policy**.
5. In the **Quotas** column, click next to the **group** or **object** for which you want to set a quota.



6. In the **Quota Policy** window enter the amount of storage space you want to allocate for a quota.



7. Click **Save** to set the amount.

Quotas and Low Space Errors

Be very careful setting quotas. As a group or virtual database (VDB) approaches the quota level, snapshots may fail and logs may not be captured, causing LogSync to fail.

Quota Thresholds

The following is a table of the thresholds and a description of the actions taken at each threshold. This behavior is generic across **all** engines.

Thresholds	Default Ranges	Actions Taken
Critical	Quota \geq 95%	<ul style="list-style-type: none"> • Disable VDBs and dSources under quota (for a group quota, everything inside gets disabled). • Disallow jobs detailed in the DB actions table below. • Cancel all in-flight jobs listed in the DB actions table except for enabling. • Poll every minute.
Resume	95% > Quota \geq 90%	<ul style="list-style-type: none"> • VDBs and dSources disabled by critical quotas will not be automatically re-enabled (DBs can still be manually re-enabled). • Resume fault is triggered if the quota has fallen from the critical threshold and if a DB was disabled from hitting the critical threshold. • Warning fault is triggered. • Poll every 3 minutes.
Warning	90% > Quota \geq 80%	<ul style="list-style-type: none"> • Warning fault is triggered. • VDBs and dSources disabled from hitting critical quota will be re-enabled. • Poll every 3 minutes.
Safe	80% > Quota	<ul style="list-style-type: none"> • VDBs and dSources disabled from hitting critical quota will be re-enabled. • Poll every 5 minutes.

If a DB was disabled prior to reaching the **critical** quota, it will not be automatically re-enabled when falling to an acceptable range or when the quota is removed.

Disallowed Database Actions when in Quota Critical Threshold

Here is a table of the actions that are not allowed when the quota is in the critical threshold.

DB Actions Table

Target is Group	Target is not Group
<ul style="list-style-type: none"> • Refresh target in the group. • Sync target in the group. • Link into the group. • Provision into the group. • Enable target in the group. 	<ul style="list-style-type: none"> • Refresh target in the group. • Sync target in the group. • Enable target in the group.

If a **critical** quota is placed on a target, dSources, and VDBs can still provision to groups that are quota below the **critical** threshold.

Related Topics

- [Capacity and Resource Management](#)

Deleting Objects to Increase Capacity

This topic describes how to delete database objects to create additional capacity.

Deleting unused or outdated objects should be a regular part of Delphix Engine administration. This is especially important to prevent low space errors, which can cause the Delphix Engine to stop. The Delphix Engine holds a maximum of 400 objects.

1. Log into the **Delphix Management** application.
2. Select **Resources > Storage Capacity**.
3. Next to the object you want to delete select the **Trashcan**.
4. In the **Delete** dialog, select **Force Delete**. Oracle users will have the option to provide additional credentials.



The screenshot shows a dialog box titled "Delete Dataset Child VDB" with a close button (X) in the top right corner. The main text asks, "Are you sure you want to delete dataset 'Child VDB'?". Below this text are two checkboxes: "Force Delete" and "Provide privileged credentials". At the bottom right of the dialog are two buttons: "Cancel" and "Delete".

5. Click Delete.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot. These items are displayed with a lock icon next to the name.

Delete Dataset dbdhcp1 ✕

Unable to delete dataset dbdhcp1

Dataset dbdhcp1 is locked due to the following dependencies:

- VDB "C3" has been provisioned from it
- VDB "C1" has been provisioned from it

Dataset C3 is locked due to the following dependencies:

- VDB "C4" has been provisioned from it
- VDB "C6" has been provisioned from it
- Self Service template "JSDataTemplate(C3)" has a reference to it

Dataset C4 is locked due to the following dependencies:

- VDB "C5" has been provisioned from it
- Self Service container "JSContainer(C4)" has a reference to it

Dataset C1 is locked due to the following dependencies:

- VDB "C2" has been provisioned from it

[Copy to Clipboard](#) [Close](#)

Related Topics

- [Capacity and Resource Management](#)
- [Understanding Delphix Disk Usage](#)

Adding and Expanding Storage Devices

Prerequisites

If you are expanding a storage device after initial configuration, first make sure to add capacity to it using the storage management tools available through the device's operating system. In vSphere, for example, you can add capacity using **Edit System Settings**.

VMware Hypervisor (vSphere)

It is usually not necessary to reboot the Delphix Engine to add/expand storage if you're using vSphere. If you followed the guidelines for adding your initial storage using all 4 virtual SCSI controllers, the Delphix Engine will see new storage without a reboot. On the other hand, if your new storage is leveraging a new virtual SCSI controller, you will need to reboot the Delphix Engine to detect the new storage. See [Deployment for VMware](#) for more information.

Adding a Storage Device

If it is available you can add more storage devices to the Delphix Engine.

1. Launch the **Delphix Setup** application and log in using the **sysadmin** credentials.
2. In the **Storage** section of the **Server Setup Summary** screen, click **Modify**.
3. The Delphix engine should automatically detect any new storage devices. If a newly added storage device does not appear in the **Storage** section of the **Server Setup Summary** screen, click **Rediscover**.
4. The newly-added devices will have a drop-down in the **Usage Assignment** column. Set the **Usage Assignment** to **DATA** for newly-added devices that you wish to add to the storage pool.
5. Click **Save**.

Expanding a Storage Device

1. Launch the **Delphix Setup** application and log in using the **sysadmin** credentials.
2. In the **Storage** section of the **Server Setup Summary** screen, click **Modify**.
3. Select **Expand** for each device that you want to expand. The **Expand** checkbox appears next to the name of devices that have added capacity (in other words, the underlying LUN has been expanded), and the **Unused** column indicates how much capacity is available for each device.
4. Click **Save**.

Related Topics

- [Capacity and Resource Management](#)
- [Deployment for VMware](#)

Delphix Storage Migration

⚠ Multiple Device Removal in the Delphix Engine 6.0.12.0 and higher version introduces a breaking kernel module change that requires a reboot to load the new module. Therefore, a deferred reboot engine upgrade operation will be unable to remove devices until a reboot is performed.

Getting Started

Delphix storage migration allows you to remove storage devices from your Delphix Engine, provided there is sufficient unused space, thus allowing you to repurpose the removed storage device to a different server. You can also use this feature to migrate your Delphix Engine to different storage by adding the new storage devices and then removing the old storage devices.

⚠ Feature Compatibility

This feature is only compatible with Delphix Engine Releases 5.0.4 and later.

Possible Migration Methods

Method	Pros	Cons
Delphix Storage Migration	<ul style="list-style-type: none"> • Good for migrating storage that was accidentally added to the engine or added to the engine improperly (wrong size). • May reduce fragmentation if new storage is added to replace old disks. • Good for migrating a small amount of storage (e.g.: < 10 TB). 	<ul style="list-style-type: none"> • With Delphix versions prior to 5.3, this mapping table can consume 2-3GB of RAM for every 1TB of allocated data that is migrated, if the disk being removed has a high level of fragmentation. From version 5.3, DxFs will migrate the data in larger blocks, comprising both allocated and unallocated space. This allows for significantly fewer mapping entries, with memory usage typically reduced to 50-100MB per TB of allocated data that is migrated. • May increase fragmentation on remaining disks if no new disks are added. • Depending on the size of the disk and storage performance this method could be less performant than other methods. • Each device removed could take longer than the previous one as data is remapped across the remaining disks. • The maximum number of 20 devices can be removed in releases prior to Delphix version 5.1.

Method	Pros	Cons
vMotion	<ul style="list-style-type: none"> Fast 	<ul style="list-style-type: none"> If there is high fragmentation on the existing disks, this is copied to the new disks.
Delphix Replication	<ul style="list-style-type: none"> Data is completely rewritten from one Delphix Engine to another which significantly reduces fragmentation on the new Delphix Engine. Replication can be configured to limit the impact on the network (compression and bandwidth). Replication is resumable from network disruptions, on in the event of replication source or target stack/host restarts. It is currently NOT possible to manually suspend/resume a replication job. 	<ul style="list-style-type: none"> Depending on the number of objects to replicate as well as network and storage performance, this method could be considered slow. This does require an outage to "migrate" the objects from the replication source Delphix Engine to the replication target Delphix Engine - outage time depends on several factors like the number of objects, incremental replication time, time to enable/disable objects, etc. Only migrates storage objects like VDBs/dSources, and dependent environment information. Other items like users/policies/events/job history/config templates are NOT replicated.

Understanding Delphix Storage Migration

Delphix storage migration is a multi-step process:

1. Identify a storage device for removal. The device you choose will depend on your use case.
 - a. To remove extra storage that is unused, you can select any device for removal. For best performance, select the device with the least allocated space; typically, this is the device that you added most recently. The allocated space is defined by the `usedSize` property of the storage device:

```
test-env@delphix 'Disk10:2'> ls
Properties
  type: ConfiguredStorageDevice
  name: Disk10:2
  bootDevice: false
  configured: true
  expandableSize: 0B
  model: Virtual disk
  reference: STORAGE_DEVICE-6000c293733774b7bb0e4aea83513b36
  serial: 6000c293733774b7bb0e4aea83513b36
  size: 8GB
  usedSize: 7.56MB
  vendor: VMware
```

- b. To migrate the Delphix Engine to new storage, add storage devices backed by the new storage to the Delphix Engine. Then remove all the devices on the old storage.
2. Use the Delphix command-line interface (CLI) to initiate the removal of your selected device.

3. Data will be migrated from the selected storage device to the other configured storage devices. This process will take longer the more data there is to move; for very large disks, it could potentially take hours. You can cancel this step if necessary.
4. The status of the device changes from **configured** to **unconfigured** and an alert is generated to inform you that you can safely detach the storage device from the Delphix Engine. After this point, it is not possible to undo the removal, although it is possible to add the storage device back to the Delphix Engine.
5. Use the hypervisor to detach the storage device from the Delphix Engine. After this point, the Delphix Engine is no longer using the storage device, and you can safely re-use or destroy it.

Limitations of Delphix Storage Migration

After removal, the Delphix Engine uses memory to track the removed data. In the worst-case scenario, this could be as much as 2-3 GB of memory per TB of used storage. Note that this is used storage; the overhead of removing a 1TB device with only 500MB of data on it will be much lower than the overhead of removing a 10GB device with 5GB of data on it.

User Interface

Delphix storage migration is currently available exclusively via the CLI. There are two entry points.

- `verifyStorageDeviceRemoval` - Verifies available pool space and available memory.
- `startStorageDeviceRemoval` - Kicks off job, detailed discussion of job concurrency/recovery.

Device Removal for Storage Migration

DO NOT REMOVE A CONFIGURED STORAGE DEVICE

Do not remove a configured storage device or reduce its capacity. Removing or reducing a configured storage device will cause a fault with the Delphix Engine, and will require the assistance of Delphix Support for recovery.

1. Identify which device you want to remove.
 1. If you are using a VMware **RDM** disk, note the UUID of the device by looking at its name in the vSphere GUI. For more information, see the [Getting the UUID of a RDM Disk from VMware, via the vSphere GUI](#) article.
 2. If you are using a VMware **virtual** disk, note the UUID of the device via the vSphere API. See the section of this [VMware KB article](#) on how to get the UUID of your virtual disk.
 3. In EC2, note the attachment point – e.g., `/dev/sdf`.
 4. In KVM, note the UUID.
 - Login to the Delphix CLI as a **sysadmin** user.
 - Navigate to the storage/device directory with `cd storage/device`.
 - Select one or more devices to remove and make note of the name (e.g., select "Disk10:2" and "Disk10:3").

```
test-env@delphix storage/device> ls
Objects
NAME      CONFIGURED  SIZE  EXPANDABLESIZE
Disk10:2  true        8GB   0B
Disk10:0  true        24GB  0B
Disk10:1  true        8GB   0B
Disk10:3  true        8GB   0B
```

- (For VMware only) Confirm the disk selection is correct by validating that the serial matches your UUID:

```
test-env@delphix storage device 'Disk10:2'> ls
Properties
  type: ConfiguredStorageDevice
  name: Disk10:2
  bootDevice: false
  configured: true
  expandableSize: 0B
  model: Virtual disk
  reference: STORAGE_DEVICE-6000c2909ccd9d3e4b5d62d733c5112f
  serial: 6000c2909ccd9d3e4b5d62d733c5112f
  size: 8GB
  usedSize: 8.02MB
  vendor: VMware
```

- Navigate to `storage/remove`. Run `verify` and set the devices to be removed, as shown below.

```
test-env@delphix storage remove > verify
test-env@delphix storage remove verify *> set devices="Disk10:2,Disk10:3"
test-env@delphix storage remove verify *> commit
  type: StorageDeviceRemovalVerifyResult
  newFreeBytes: 15.85GB
  newMappingMemory: 3.14KB
  oldFreeBytes: 23.79GB
  oldMappingMemory: 0B
```

- Navigate to `storage/remove`. Run `start` and set the devices to be removed, as shown below.

```
test-env@delphix storage remove > start
test-env@delphix storage remove start *> set devices="Disk10:2,Disk10:3"
test-env@delphix storage remove start *> commit
  Dispatched job JOB-4
  STORAGE_DEVICES_START_REMOVAL job started.
  STORAGE_DEVICES_START_REMOVAL job completed successfully.
```

 This does not signify that the device migration has been completed. A `STORAGE_DEVICE_REMOVAL` job will start for each device to be removed, which handles the data migration from that disk.

- Wait for device evacuation to complete. Alternatively, you can cancel the evacuation.

! Do not detach the device from the Delphix Engine in your hypervisor until the data evacuation is completed.

You can monitor the progress of the `STORAGE_DEVICE_REMOVAL` job in the Management GUI under **System > Jobs**.

- Once the device evacuation has been completed, the job will finish and a fault will be generated. Detach the disks from your hypervisor and the fault will clear on its own. An example of the fault created is shown below.

The screenshot shows the 'DELPHIX MANAGEMENT' interface with a 'Faults' section. The 'Current' tab is active, showing a table with one fault entry:

Severity	Diagnosed	Title	Target
WARNING	Nov 2, 2021 1:54 PM	Device evacuation completed	Disk2:2

To the right of the table is a detailed view of the selected fault, showing a 'WARNING' icon, the date 'Nov 2, 2021 1:54 PM', the title 'Device evacuation completed', the target 'Disk2:2', and details: 'The evacuation of Disk2:2 (acsi-36000c29da4d7af6c0d5d4bf9d38784ba03) has been completed.' The user action is 'Remove the device from the hypervisor.' At the bottom right of the details panel are 'Resolve' and 'Ignore' buttons.

i Using VMDKs

When using VMDKs, deleting the wrong VMDK could cause data loss. Therefore, it is highly advisable to detach the device, then verify that the Delphix Engine continues to operate correctly, and lastly delete the VMDK.


Getting the UUID of an RDM Disk from VMware via the vSphere GUI.

! In the event that the disk serial number displayed in Delphix does not match the UUID in VMware, the Delphix Engine must be powered off and back on in order to make VMware provide the correct values to the guest operating system (Delphix). This has been necessary when using **vmkfstools** with **setuuid**. When forcing the guest OS to re-read the SCSI sense data for the device, VMware still provides the original values. Even after a simple reboot VMware still provides the previous UUID values. It was not until the VM was explicitly powered off and back on did VMware present the new UUID values to the guest. After which the UUIDs matched between the **vmkfstools getuuid** command and the CLI output.


In the ESX graphical user interface (GUI), select your **VM**.

1. Click **Edit settings**.
2. If not already displayed, select the **Hardware** tab.
3. Select the **device** you want to remove.

4. Click **Manage Paths**.



The UUID of the device appears in the title bar, as seen below.



Getting the UUID of a VMDK from VMware, via ssh to the ESX server

1. ssh onto the ESX server as the root user.
2. Navigate to the directory containing the .vmdk files for the Delphix VM.

Use the 'vmkfstools -J getuuid <.vmdk filename>' command to obtain the UUID, for example:

```
/vmfs/volumes/25894daa-f7b2b044/delphix01-2356 # vmkfstools -J getuuid
delphix01-2356_1.vmdk
UUID is 60 00 C2 91 01 bc 8e 72-31 a4 cd b0 b3 f6 e5 74
```

Getting the UUID of a VMDK from VMware, via VMware PowerCLI

```
PS C:\> Connect-VIServer -Server durban -Protocol https -Username root
-PassWord root_password
```

Name	Port	User
----	----	----
durban	443	root

```
PS C:\> Get-VM delphixVM | Get-HardDisk | select
name,filename,@{name="UUID";expr={$_.extensiondata.backing.uuid}}
```

Name	Filename	UUID
----	-----	-----
Hard disk 1		[zfs_delphixVM]
d\px-5.1.9.0-432-61155cf...	6000C294-a115-0327-e417-02560d86e944	
Hard disk 2		[zfs_delphixVM]
d\px-5.1.9.0-432-61155cf...	6000C299-38fe-5050-1eb2-1ee6db62b257	
Hard disk 3		[zfs_delphixVM]
d\px-5.1.9.0-432-61155cf...	6000C294-662d-c674-8957-03e0514b7006	
Hard disk 4		[zfs_delphixVM]
d\px-5.1.9.0-432-61155cf...	6000C29d-0719-1072-0f85-96da2efef4a3	

```
PS C:\> Disconnect-VIServer
```

Confirm

Are you sure you want to perform this action?

Performing operation "Disconnect VIServer" on Target "User: root, Server: durban, Port: 443".

[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
(default is "Y"): Y

Related Topics

- [Capacity and Resource Management](#)
- [Adding and Expanding Storage Devices](#)


Managing Source Data

The **Managed Source Data** dashboard tile provides an aggregate sum of the size of ingested source databases. This helps to understand the total amount of data per-engine that is managed by Delphix.

Managed Source Data refers to the size of the source data that is ingested and managed by the Delphix Engine. It is the physical, allocated size of the source database.

Today, the value displayed represents the source data size for Oracle, SQL Server, and SAP ASE databases. Since this value is a measure of source size, it is not affected by Delphix constructs or operations, e.g. the sizes of dSources, VDBs, snapshots. Since the intent is to measure the size of the sources, data from replication or Selective Data Distribution are not included in the sum for the engine. The total usage is displayed in gigabytes.

When an admin (Delphix Management Application) or sysadmin (Delphix Setup) login to the Delphix Engine they will have access to the Managed Source Data tile (shown below) from their dashboard.



Selecting the **View** link located on the top-right takes users to the Managed Source Data page. This page provides a detailed breakdown of the total usage and can also be reached via **Resources > Managed Source Data** from the navigation bar.

Blue question marks indicate:

- A “fallback” value is being used in the sum: A fallback value is used when sources are unavailable at the time when the usage query is made (e.g., the source is disabled), a “fallback” value is used in the sum. The fallback value refers to the last usage value Delphix collected while connected to the source.
- Sources that are not included in the sum: These are source types for which Delphix does not currently collect data (for example, sources like EBS, DB2, HANA, PostgreSQL).

Hovering over the **blue question marks**, a tooltip appears to describe the exact situation (whether or not Delphix used historical data/fallback values or if the source for collection is unsupported).

Managed Source Data

Managed Sources How It Works

Filter: none  

Name	Last collected	Space used	Type Of Data
CDOMSHSRD97A	Sep 29, 2021 12:24 PM	2.22 GB	Oracle
DBOMSRBDC6C	Sep 29, 2021 12:24 PM	2.12 GB	Oracle
CDOMLOTG4F5E	Sep 20, 2021 2:00 PM	1.69 GB 	Oracle
CDOMSHSRD97APDB1	Sep 29, 2021 12:24 PM	773.63 MB	Oracle
CDOMSHSRD97APDB2	Sep 29, 2021 12:24 PM	763.63 MB	Oracle
CDOMSHSRD97APDB3	Sep 29, 2021 12:24 PM	763.63 MB	Oracle
CDOMLOTG4F5EPDB3	Sep 20, 2021 2:00 PM	762.84 MB 	Oracle
Delphix_Admin	Sep 17, 2021 10:30 PM	3.88 MB 	MSSQL Server
CDOMLOSRSB5A3PDB1	Sep 29, 2021 12:24 PM	Copy query 	Oracle

Total Space Used

9.03

GB*

Sep 29, 2021

MSSQL Server	3.88 MB
Oracle	9.02 GB

*Space usage data is available when Delphix sources are able to be monitored.

To search the **Managed Source Data** table, enter the name of the data source you are looking for. The grid will refresh to display the selected data source. Select  to refresh data to get updated space usage information. To export the information provided in the grid to a .csv file select .

 **Backend data refresh**

The backend data refreshes at a regular interval but does not trigger a UI refresh.

Related Topics

- [Capacity and Resource Management](#)

An Overview of Held Space

Scenario Description

The existence of a Held Space is not indicative of an issue but rather a representation of an underlying storage dependency.


There are three scenarios that can create Held Space.

Scenario 1:

The first scenario can occur when you replicate a dSource to a second engine, which we'll call a Replication Target.

If you then provision a VDB from the Replication Target, and the dSource is subsequently deleted from the Replication Source, you will create a Held Space when that delete replicates to the Replication Target engine.

When this happens, the deleted dSource is removed from the target, but its storage remains held because it is needed by the replica provisioned VDB.



Held Space can be viewed in the **Storage Capacity** screen by selecting the Held Space tab.

Object	Size
STORAGE_CONTAINER-96	74.49MB

Scenario 2:

The second scenario can occur in the context of a snapshot that is in use by a Delphix Self-Service branch or bookmark.

The snapshot may be removed via policy, but until the branches and/or bookmarks that leverage that snapshot are removed, its space will be held.

Scenario 3:

The third scenario can occur during Selective Data Distribution where a masked VDB that is provisioned from a dSource is replicated to a Replication Target.

The metadata objects relating to the dSource could contain sensitive information like hostnames, user names, and passwords. Hence they are not replicated to the Replication Target. The data belonging to the dSource is replicated after redacting any sensitive information. This redacted data shows up as held space.

Scenario 4:

The first scenario is caused by a failed or interrupted Replication job and results in held space on the Replication source engine. During replication, all data blocks to be replicated are serialized for transfer. Once the transfer is complete, the serialized data is released.

However, if a Replication job fails, the serialized data is retained until the next successful Replication job execution, or the Replication Profile is deleted - whichever happens first.

In a scenario where dSource and VDB snapshots in the serialized data reach the Retention Policy limit before the Replication job successfully executes, that snapshot data will become held space. This is most likely to happen if either of the following is true:

- A failed Replication job goes a long time before being successfully re-run.
- Replication schedules are configured with a period that is longer than one or more Retention Policies.

Related Topics

- [Delphix Storage Migration](#)
- [Capacity and Resource Management](#)

Monitoring and Log Management

Using both Delphix and its associated datasets will generate many types of logs. Monitoring these logs are a key component of good diagnosability. Here, we explain the types of logs that Delphix creates as well as the monitoring tool integrations we support, such as SNMP and Splunk.

This section covers the following topics:

- [Configuring SNMP](#)
- [Viewing Action Status](#)
- [Viewing Jobs](#)
- [System Faults](#)
- [Accessing Audit Logs](#)
- [Creating Support Logs](#)
- [Setting Support Access Control](#)
- [Setting Syslog Preferences](#)
- [Diagnosing Connectivity Errors](#)
- [Email \(SMTP\) Alert Notifications](#)
- [Splunk Integration](#)
- [Fluentd Plugin Service for API Modules](#)

Configuring SNMP

This topic describes how to configure SNMP.

Starting with version 6.0.6.0 the Delphix Engine adds support for SNMP version 3 by implementing the User Security Model (USM) of the SNMP version 3 (SNMPv3) specification. Version 3 of SNMP adds stronger security compared to version 2.

 When SNMP v2 is configured and enabled, SNMP v1 is also configured and enabled. Similarly, v2 is configured and enabled with v3. This is a known behavior of the net-snmp package used to implement SNMP on the engine.

USM replaces the community string as in SNMP version 2 with user records. The Delphix engine can both receive and send SNMP version 3 messages only by using the USM security model. The USM model provides more security over version 2 by hashing passwords and encrypting the payload.

The Delphix Engine includes an SNMPv3 agent that is only capable of responding to read-only operations, such as GET messages using the User Security Model. USM users can be configured via the GUI, CLI, or API. Each user is composed of a username, authentication password, authentication protocol, privacy password, and a privacy protocol.

The Delphix Engine is also capable of sending alerts and faults over SNMP with the User Security Model. Messages are sent identically as with SNMP version 2 except that the USM model is used to authenticate and encrypt the alert message against a customer's manager. Users will need to properly configure their managers to receive messages. The use of the test function can be used to test configuration.

Prerequisites

There are no prerequisites for enabling SNMP to provide system status. The following are prerequisites for sending alerts to an external SNMP manager.


- At least one SNMP manager must be available and must be configured to accept SNMPv2 InformRequest notifications.
- Delphix's MIB (Management Information Base) files must be installed on the SNMP manager or managers. These MIB files describe the information that the Delphix Engine will send out. They are attached to this topic.

File Name	Content-Type	File Size
DELPHIX-ALERT-MIB.txt	text/plain	5 kB
DELPHIX-MIB.txt	text/plain	0.8 kB

Configuring SNMP for v2

1. On the Delphix Engine login screen, select **Delphix Setup**.
2. In the Delphix Setup login screen, enter the **sysadmin** username and password.
3. Click **Log In**.

4. From the **Delphix Setup**, select **Preferences > SNMP Configuration**.



5. Select **Enable**.
6. Select **SNMP Version 2**.

 A screenshot of the 'SNMP Configuration' dialog box. The title bar says 'SNMP Configuration' with a close button (X) on the right. The main content area is divided into sections:

- Simple Network Management Protocol**: A checkbox labeled 'Enable' is checked.
- SNMP Version**: Two radio buttons are present: 'SNMP Version 3 (recommended)' (unselected) and 'SNMP Version 2' (selected).
- System Status** and **Warnings and Alerts**: Two tabs are shown, with 'System Status' selected. Below them is the text 'Settings for SNMP GET request to query system status.'
- Community String**: A text input field is empty.
- Authorized Network**: Three radio buttons are present: '0.0.0.0/0 (Allow all clients)' (selected), '127.0.0.1/32 (Block all clients)' (unselected), and 'Custom' (unselected) with an adjacent text input field.
- System Location**: A text input field contains the word 'Unknown'.


 At the bottom right, there are 'Cancel' and 'Save' buttons.

7. In **Community String**, enter the community string. It is the string that SNMP clients must provide in order to be authorized to retrieve SNMP information from the Delphix Engine.
8. OPTIONAL: change the following settings:
- Authorized Network** – The set of client IP addresses (in CIDR notation) authorized to retrieve SNMP information from the Delphix Engine. To allow all clients, set this to 0.0.0.0/0 (the default). To prevent all clients from connecting, set this to the loopback address, 127.0.0.1/32.
 - System Location** – A free-form text description of the Delphix Engine's physical location. This is provided as the value for MIB-II OID.

9. Click **Save** to commit the SNMP configuration.

Configuring SNMP for v3

1. On the Delphix Engine login screen, select **Delphix Setup**.
2. In the Delphix Setup login screen, enter the **sysadmin** username and password.
3. Click **Log In**.
4. From the **Delphix Setup**, select **Preferences > SNMP Configuration**.



5. Select **Enable**.
6. Select **SNMP Version 3** (recommended).
7. Click **Save** to commit the SNMP configuration.

Configure SNMP Version 3 Agent

After enabling SNMP and choosing version 3, the SNMP agent can be configured with the User Security Model (USM) users.

1. Click on **System Status**.

2. Click on the + to show the **New SNMP USM User Information** dialog.


The screenshot shows a dialog box titled "New SNMP USM User Information". It contains the following fields and values:

- Username:** John Smith
- Authentication Method:** SHA
- Authentication Passphrase:** (masked with asterisks)
- Encryption Method:** AES
- Encryption Passphrase:** (masked with asterisks)

At the bottom right of the dialog are two buttons: "Cancel" and "Save".

3. Provide the following information for the user:
- Username** - A string representing the name of the user.
 - Authentication Method** - An indication of whether messages sent on behalf of this user can be authenticated, and if so, the type of authentication protocol that is used. Users can select from:
 - MD5** -the HMAC-MD5-96 authentication protocol.
 - SHA** -the HMAC-SHA-96 authentication protocol.
 - Authentication Passphrase** - If messages sent on behalf of this user can be authenticated, this key will be used with the authentication key for use with the authentication. This field requires a minimum of 8 characters.
 - Encryption Method** - Provides support for data confidentiality. The designated portion of an SNMP message is encrypted and included as part of the message sent to the recipient. Users can select from
 - AES** - Advanced Encryption Standard
 - DES** - Data Encryption Standard
 - Encryption Passphrase** - If messages sent on behalf of this user can be encrypted, this key will be used with the encryption method for use with the authentication. This field requires a minimum of 8 characters.

4. Click **Save** to add your user.



The image shows a 'SNMP Configuration' dialog box with a close button (X) in the top right corner. The dialog is titled 'SNMP Configuration' and contains the following sections:

- Simple Network Management Protocol**: A checkbox labeled 'Enable' is checked.
- SNMP Version**: Two radio buttons are present: 'SNMP Version 3 (recommended)' (selected) and 'SNMP Version 2'.
- System Status** and **Warnings and Alerts**: Two tabs are visible, with 'System Status' selected and underlined.
- Settings for SNMP GET request to query system status.**: A heading for the current tab.
- USM Users**: A table with a '+' icon on the right. The table has four columns: 'Username', 'Authentication', 'Encryption', and 'Actions'. One row is visible with the following data:

Username	Authentication	Encryption	Actions
John Smith	SHA	AES	 
- System Location**: A text input field containing the value 'Unknown'.
- Buttons**: 'Cancel' and 'Save' buttons are located at the bottom right of the dialog.

Configuring SNMP Version 3 Managers

1. Click on the **Warnings and Alerts** tab.
2. Choose the **Severity** of the alerts that are sent via SNMP. The severity is ranked in order from most restrictive to least restrictive: critical, warning, and informational. Setting the severity critical will only include critical alerts while setting to informational will cause an informational, warning and critical alerts to be sent over SNMP to all managers.

3. Click **+** to add a manager to show the **New SNMP Manager Information** dialog.

New SNMP Manager Information ✕

Username

Host Address

Port

Authentication Method

None ▾

Authentication Passphrase

Encryption Method

None ▾

Encryption Passphrase

Use INFORM instead of TRAP.

Cancel Save

4. Provide the following information:
 - a. **Username** - A string representing the name of the user.
 - b. **Host Address** - Host or IP address.
 - c. **Port** - Port used by the SNMP Engine.
 - d. **Authentication Method** - An indication of whether messages sent on behalf of this user can be authenticated, and if so, the type of authentication protocol that is used. Users can select from:
 - MD5** -the HMAC-MD5-96 authentication protocol.
 - SHA** -the HMAC-SHA-96 authentication protocol.
 - e. **Authentication Passphrase** - If messages sent on behalf of this user can be authenticated, this key will be used with the authentication key for use with the authentication. This field requires a minimum of 8 characters.
 - f. **Encryption Method** - Provides support for data confidentiality. The designated portion of an SNMP message is encrypted and included as part of the message sent to the recipient. Users can select from
 - AES** - Advanced Encryption Standard
 - DES** - Data Encryption Standard
 - g. **Encryption Passphrase** - If messages sent on behalf of this user can be encrypted, this key will be used with the encryption method for use with the authentication. This field requires a minimum of 8 characters.
5. By default, TRAP messages are sent and require the agent's authoritative engine id for the user to be configured with the SNMP engine ID from the Delphix Engine, which can be obtained via the CLI. Click Use **Inform instead of Trap** to use INFORM messages, which are more reliable and require less configuration at the expense of higher network traffic.
6. Click **Save**.

CLI: Viewing SNMP Engine ID

Viewing SNMP Engine ID

```
> service snmp v3
service snmp v3> ls
Properties
  type: SNMPV3Config
  enabled: true
  engineId: 0x80001f88801ad80b5e62a7f85f0000000
  location: Unknown
  securityModel: USM
  severity: WARNING

Children
manager
usm

Operations
update
```

Supported MIBs

The Delphix software can be configured to send SNMP traps when Delphix alerts are generated as described in the procedure above. In order to process these traps in your SNMP manager software, you will need the base Delphix MIB and the Delphix Alert MIB.

In addition to generating traps, the Delphix Engine supports read-only access to the following MIBs for basic system monitoring purposes.

- The following MIB-II object hierarchies are defined in [RFC 1213](#):
 - system (OID .1.3.6.1.2.1.1): Provides basic system identity information
 - interfaces (OID .1.3.6.1.2.1.2): Provides network interface information including I/O statistics
 - IP (OID .1.3.6.1.2.1.4): IP protocol information including IP addresses configured, routes, and IP statistics
- The following UCD objects defined in <http://www.net-snmp.org/docs/mibs/ucdavis.html>
 - memory usage (OID .1.3.6.1.4.1.2021.4)
 - CPU usage (OID .1.3.6.1.4.1.2021.11)
 - Disk I/O statistics (OID .1.3.6.1.4.1.2021.13.15)

Examples

The following examples assume that you have enabled SNMP on a Delphix Engine named example.company.com, and have set the community string to "public".

1. Walking the MIB-II objects using the net-snmp snmpwalk tool:

```
$ snmpwalk -v 2c -c public example.company.com
SNMPv2-MIB::sysDescr.0 = STRING: Delphix Engine 5.1.6.0 DelphixOS 5.1.2017.03.2
4
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.41028
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (476432) 1:19:24.32
SNMPv2-MIB::sysContact.0 = STRING: administrator@company.com
SNMPv2-MIB::sysName.0 = STRING: example.company.com
SNMPv2-MIB::sysLocation.0 = STRING: VM Host
...
```

2. Walking Disk read and write I/O statistics:

```
$ snmpwalk -v 2c -c public example.company.com .1.3.6.1.4.1.2021.13.15.1.1.12
UCD-DISKIO-MIB::diskIONReadX.1 = Counter64: 11310593921
UCD-DISKIO-MIB::diskIONReadX.2 = Counter64: 334
UCD-DISKIO-MIB::diskIONReadX.3 = Counter64: 334
UCD-DISKIO-MIB::diskIONReadX.4 = Counter64: 865912605
UCD-DISKIO-MIB::diskIONReadX.5 = Counter64: 867599133
UCD-DISKIO-MIB::diskIONReadX.6 = Counter64: 865339677
UCD-DISKIO-MIB::diskIONReadX.7 = Counter64: 11309258752
UCD-DISKIO-MIB::diskIONReadX.8 = Counter64: 0
UCD-DISKIO-MIB::diskIONReadX.9 = Counter64: 1822880256
seb-laptop:~$ snmpwalk -v 2c -c public example.company.com .1.3.6.1.4.1.2021.1
3.15.1.1.13
UCD-DISKIO-MIB::diskIONWrittenX.1 = Counter64: 22337830400
UCD-DISKIO-MIB::diskIONWrittenX.2 = Counter64: 0
UCD-DISKIO-MIB::diskIONWrittenX.3 = Counter64: 0
```

```
UCD-DISKIO-MIB::diskIONWrittenX.4 = Counter64: 45118203392
UCD-DISKIO-MIB::diskIONWrittenX.5 = Counter64: 45137660928
UCD-DISKIO-MIB::diskIONWrittenX.6 = Counter64: 45139064320
UCD-DISKIO-MIB::diskIONWrittenX.7 = Counter64: 22337830400
UCD-DISKIO-MIB::diskIONWrittenX.8 = Counter64: 0
UCD-DISKIO-MIB::diskIONWrittenX.9 = Counter64: 33023515648
```

3. Retrieving the system uptime:

```
$ snmpget -v 2c -c public example.company.com .1.3.6.1.2.1.1.3.0
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1453172) 4:02:11.72
```

Related Topics

- [Monitoring and Log Management](#)

Viewing Action Status

This topic describes how to view the status of actions for the Delphix Engine.

To view the status of actions that are currently running on the Delphix Engine, open the **Action sidebar**. To view details of currently-running and completed jobs, open the **Dashboard**.

Action Sidebar Procedure


1. Login to the **Delphix Management** or **Delphix Setup** application.
Depending on the width of the window, the **Action sidebar** may be automatically displayed on the right of the screen.
2. To see the **Action sidebar**, click **Actions** on the top navigation bar.

Description

The **Action sidebar** consists of two sections. The top section lists actions that are currently running on the Delphix Engine. The bottom section, labeled **Recently completed**, contains actions that have recently been completed.

Each action is initially collapsed and only presents the title of the action. Click an action to expand it and see more details such as progress, elapsed time, and a description of the operation in progress.

The following is an example of the **Action sidebar** when an Enable action is running.



When an action has been completed it will move down to the bottom of the panel under Recently completed. The following is an example of the **Recently completed** section when a Refresh action is has completed.

Recently completed	
Refresh and discover environment "Starfire".	
Refresh environment "Starfire".	✓
Disable a list of datasets in environment "S...	✓
Refresh hosts in environment "Starfire".	✓
Refresh host "bbdhcp-tgt-AHCI-58503.dc...	✓
Enable a list of datasets in environment "St...	✓
Discover information and objects for environ...	✓
Update repository "BBDHCP_S1".	✓
Update Oracle listener "LISTENER".	✓

If you are a Delphix Admin or a System Admin, you will be able to see all the actions of your respective application. If you are not an admin user, you will only see actions you have permission to see.

Sub-action

Each action may contain one or several sub-actions which represent the execution of a subset of the action itself. Click an action to see its sub-actions and their respective details. Note that the list of sub-actions is created dynamically during the execution of the action.

The following is an example of an Environment Refresh action and its sub-actions.

Actions	
Refresh and discover environment "Starfire".	
<div style="background-color: #ccc; width: 100%; height: 10px;"></div>	34 s
Refresh environment "Starfire".	
<div style="background-color: #ccc; width: 100%; height: 10px;"></div>	34 s
Disable a list of datasets in environment "S... ✓	
Refresh hosts in environment "Starfire".	
<div style="background-color: #007bff; width: 25%; height: 10px;"></div> <div style="background-color: #ccc; width: 75%; height: 10px;"></div>	33 s
Refresh host "bbdhcp-tgt-AHCI-58503.dc... Pushing JRE to host bbdhcp-tgt-AHCI-58503.dcenter.delphix.com [bbdhcp-tgt-AHCI-58503.dcenter.delphix.com].	
<div style="background-color: #007bff; width: 25%; height: 10px;"></div> <div style="background-color: #ccc; width: 75%; height: 10px;"></div>	33 s

Errors


When an error condition occurs during the execution of an action, the background color of the action's box becomes red, and the action remains in the top section until you dismiss it.

- Click the **action** title to expand it.
The action will expand to display a description of the error, suggestions to resolve it, and sometimes the raw output of command execution.

To dismiss the action:

- Click the **X** next to the action displaying an error.

The following is an example of an action failure displayed in the **Action sidebar**.



Related Topics

- [Monitoring and Log Management](#)

Viewing Jobs

1. Login to the **Delphix Management** application using **admin** credentials.
2. From the **System** menu select **Jobs**.

The **Jobs** panel displays all jobs that have been initiated by the Delphix Engine, and their status.

The screenshot shows the 'Jobs' panel in the Delphix Management application. The panel has a left sidebar with 'Predefined Range' (selected) and 'Custom Range' options. The main area displays a table of jobs for the period 'Dec 1, 2020 1:39:22 PM - Dec 8, 2020 1:39:22 PM'. A filter dropdown is set to 'none'. Callouts 1-4 point to the filter, the range selection, a job row, and the filter dropdown arrow respectively.

Job	Type	Description	Started	Updated	Status
JOB-443	DB_SYNC	Run SnapSync for database "vde...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:57 AM	COMPLETED
JOB-442	DB_SYNC	Run SnapSync for database "vte...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:57 AM	COMPLETED
JOB-441	DB_SYNC	Run SnapSync for database "my...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:12 AM	COMPLETED
JOB-440	DB_SYNC	Run SnapSync for database "vte...	Dec 7, 2020 7:29:05 PM	Dec 7, 2020 7:29:57 PM	COMPLETED
JOB-439	DB_REFRESH	Refresh database "vtest".	Dec 7, 2020 7:28:48 PM	Dec 7, 2020 7:29:57 PM	COMPLETED
JOB-438	SOURCE_ENABLE	Enable dataset "vtest".	Dec 7, 2020 7:24:41 PM	Dec 7, 2020 7:25:09 PM	COMPLETED
JOB-437	SOURCE_DISABLE	Disable dataset "vtest".	Dec 7, 2020 7:24:00 PM	Dec 7, 2020 7:24:29 PM	COMPLETED

1

Enter filter text to reduce the results to only those rows matching the text entered.

2


Customize results by selecting from a Predefined Range or from a Custom Range.

The screenshot shows the 'Jobs' panel in the Delphix Management application. The 'Custom Range' option is selected in the sidebar. The main area displays a table of jobs for the period 'Dec 7, 2020 1:39:22 PM - Dec 8, 2020 1:40:21 PM'. A filter dropdown is set to 'none'. Callout 3 points to the first job row.

Job	Type	Description	Started	Updated	Status
JOB-443	DB_SYNC	Run SnapSync for database "vde...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:57 AM	COMPLETED
JOB-442	DB_SYNC	Run SnapSync for database "vte...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:57 AM	COMPLETED
JOB-441	DB_SYNC	Run SnapSync for database "my...	Dec 8, 2020 3:30:00 AM	Dec 8, 2020 3:30:12 AM	COMPLETED
JOB-440	DB_SYNC	Run SnapSync for database "vte...	Dec 7, 2020 7:29:05 PM	Dec 7, 2020 7:29:57 PM	COMPLETED
JOB-439	DB_REFRESH	Refresh database "vtest".	Dec 7, 2020 7:28:48 PM	Dec 7, 2020 7:29:57 PM	COMPLETED

3

Select a job to view more details.



Timestamp	Description
Dec 7, 2020 7:25:0...	SOURCE_ENABLE job for "vtest" completed successfully.
Dec 7, 2020 7:25:0...	Dataset "vtest" enabled.
Dec 7, 2020 7:24:4...	Mounting datasets.
Dec 7, 2020 7:24:4...	Exporting storage containers from the Delphix Engine.
Dec 7, 2020 7:24:4...	Enabling dataset "vtest".
Dec 7, 2020 7:24:4...	SOURCE_ENABLE job started for "vtest".

4

Click the **Export** button to download the current page of results to a file of comma-separated values (CSV).

Related Topics

- [Monitoring and Log Management](#)

System Faults

Overview

The Faults screen provides information on states and configurations that may negatively impact the functionality of the Delphix Engine and which can only be resolved through active user intervention. When you login to the **Delphix Management** application as admin, the number of outstanding system faults appears on the right-hand side of the navigation bar at the top of the screen. Faults serve as a record of all issues impacting the Delphix Engine and can never be deleted. However, ignored and resolved faults are not displayed in the faults list.

The **Faults** screen as shown below has two tabs, **Current** and **Archive**.

Severity	Diagnosed	Title	Target
WARNING	Dec 19, 2018 2:32 AM	TCP slot table entries below recommended minimum	bbdhcp-AHCI-58503.dcenter.delphix.com
WARNING	Dec 19, 2018 2:32 AM	TCP WMEM default is below the recommended value	bbdhcp-AHCI-58503.dcenter.delphix.com
WARNING	Dec 19, 2018 2:32 AM	TCP RMEM default is below the recommended value	bbdhcp-AHCI-58503.dcenter.delphix.com
WARNING	Dec 18, 2018 2:23 PM	Invalid database credentials	Boomerang
WARNING	Dec 18, 2018 9:52 AM	TCP slot table entries below recommended minimum	bbdhcp-tgt-AHCI-58503.dcenter.delphix.com
WARNING	Dec 18, 2018 9:52 AM	TCP WMEM default is below the recommended value	bbdhcp-tgt-AHCI-58503.dcenter.delphix.com
WARNING	Dec 18, 2018 9:52 AM	TCP RMEM default is below the recommended value	bbdhcp-tgt-AHCI-58503.dcenter.delphix.com
WARNING	Dec 18, 2018 9:40 AM	Incorrect toolkit owner	bbdhcp-AHCI-58503.dcenter.delphix.com
WARNING	Dec 18, 2018 9:39 AM	Incorrect toolkit owner	bbdhcp-tgt-AHCI-58503.dcenter.delphix.com

System Faults screen

1

The number of system faults.

2

The **Faults** screen has two tabs, **Current** and **Archive**. Details of the selected fault are displayed on a card located to the right of the fault list. In the **Archive** tab, you can switch between **Resolved** or **Ignored** faults and reset all ignored faults.

3

Selecting **Refresh (Manual)**, will refresh the faults table manually, or you can select one of the other available options from the drop-down menu, available options include; Manual, 1 Second, 1 Minute, 5 Minutes.

(Note: as there is no longer a Refresh button on the screen, you must select Refresh (Manual) to refresh the screen.

4

To search the **Faults** table, enter the name of the object you are looking for. The grid will refresh to display the selected object. You can also sort using the column headings.


Resolve All will resolve all the faults in your system.

Select  to expand or close the objects in the grid.

Select  to export the information provided in the grid to a.csv file.

5

You can select and resolve multiple faults; the card panel will display how many of each type are selected. For example:



6

Details for the selected fault are displayed on a card. You can resolve or ignore faults by selecting the appropriate link at the bottom of the card.

Resolving and Ignoring Faults

Ignoring a fault will also ignore future faults of that exact type against the same object, so that future fault conditions will not be re-diagnosed even if the fault condition persists or recurs. No further notifications will be received for that specific fault condition. It is advisable to only ignore faults when the following criteria are met:

- The fault is caused by a well-understood issue that cannot be changed.
- Its impact on the Delphix Engine is well understood and does not require action.

For example, if you think that knowing about this error in the future will be important, use "Resolve" rather than "Ignore". If you reset ignored items, this clears all ignored faults, but it leaves them as resolved and does not restore the actual fault. For reset ignored faults, new faults against the same object will no longer be ignored and you will again receive notifications. Examples: If you ignore a fault "Unable to ping host" for target "192.168.1.1", Delphix ignores "Unable to ping host" errors against target 192.168.1.1. You will never see the "Unable to ping host" fault again for that target 192.168.1.1 unless you reset ignored items. Similarly, some faults are raised against snapshots which are part of a dSource. Ignoring those errors only ignores similar errors for that exact snapshot. Tomorrow's snapshot could produce the fault again.

Delphix Object-Based Environment Monitor Faults

Delphix now has a self-contained Java-based discovery infrastructure that consolidates with environmental monitoring, communicates via a common framework, and is able to provide feedback.

The environment monitor previously only created faults for "hosts" and "sources." There are several faults that more logically apply to other Delphix objects, such as repositories, which are DB install files. Posting them against sources results in fault duplication. The environment monitor now posts faults against -- and re-associates the offending faults with -- the correct objects. Consequently, users see fewer errors that are easier to diagnose.

Viewing Faults

To view the list of active system faults:

1. In the top navigation bar, click **System** then **Faults**.
2. In the Faults screen, click any fault in the list to expand it and see its details.
The details for the selected Fault will be displayed in the details card located on the right.

Each fault comprises six parts:

- **Severity** – How much of an impact the fault will have on the system. A fault can have a severity of either **Warning** or **Critical**.
 - A **Warning Fault** implies that the system can continue despite the fault but may not perform optimally in all scenarios.
 - A **Critical Fault** describes an issue that breaks certain functionality and must be resolved before some or all functions of the Delphix Engine can be performed.
- **Date** – The date that the Delphix Engine diagnosed the fault.
- **Title** – A short descriptive summary of the fault
- **Target** – The object against which the fault was posted. Faults will be posted against the host for incorrect environment configurations, sources for problems with the database, and repositories for issues with the installation.
- **Details** – A detailed summary of the cause of the fault
- **User Action** – The action you can take to resolve the fault

Addressing Faults

After viewing a fault and deciding on the appropriate course of action, you can address the fault through the user interface (UI). You can mark a fault as **Ignored** or **Resolved**. If you have fixed the underlying cause of the fault, mark it as **Resolved**. Note that if the fault condition persists, it will be detected in the future and re-diagnosed. You can mark the fault as **Ignored** if it meets the following criteria:

- The fault is caused by a well-understood issue that cannot be changed
- Its impact on the Delphix Engine is well understood and acceptable

In this case, the fault will not be re-diagnosed even if the fault condition persists. You will receive no further notifications.

To address a fault follow the steps below.

1. In the top menu bar, click **Faults**.
2. In the list of faults, click a **fault date/name** to view the fault details.
3. If the fault condition has been resolved, click **Resolve**.
Note that if the fault condition persists it will be detected in the future and re-diagnosed.
4. If the fault condition describes a configuration with a well-understood impact on the Delphix Engine that cannot be changed, you can ignore the fault by clicking **Ignore**.
Note that an ignored fault will not be diagnosed again even if the underlying condition persists.

By default, when a **critical or warning fault** occurs, the Delphix Engine immediately sends an email to the Engine Administrator (**admin**). Make sure you have configured an SMTP server and defined an appropriate email address for Engine Administrator (**admin**). See [Initial Setup](#) for more information.

Critical or Warning Alert Emails

By default, emails will also be sent for **critical or warning alerts** (aka events). You can modify the default behavior by changing the alert profile with the CLI. See the [CLI Cookbook Creating Alert Profiles](#) for more information.

Fault Lifecycle Example

Below is an image of the fault card for the fault "TCP slot table entries below the recommended minimum."

 **WARNING**

Date
Dec 19, 2018 2:32 AM

Title
TCP slot table entries below recommended minimum

Target
bbdhcp-AHCI-58503.dcenter.delphix.com

Details
The TCP sunrpc.tcp_slot_table_entries property is currently set to '16' which is below the recommended minimum value of 128.

User Action
Raise the sunrpc.tcp_slot_table_entries value to at least 128.

 [Resolving and Ignoring](#) [Resolve](#) [Ignore](#)

The **Details** section of the fault explains that the sunrpc.tcp_slot_table_entries property is set to a value that is below the recommended minimum of 128. The **User Action** section instructs you to adjust the value of the sunrpc.tcp_slot_table_entries property upward to the recommended minimum. The process for adjusting this property differs between operating systems. To resolve the underlying issue, search "how to adjust sunrpc.tcp_slot_table_entries" using a search engine and find that the second result is a [link](#) to the Delphix community forum describing how to resolve this issue. After following the instructions applicable to your operating system, return to the Delphix UI and mark the fault **Resolved**.

Related Topics

- [Monitoring and Log Management](#)
- [Viewing System Faults](#)
- [CLI Cookbook: Creating Alert Profiles](#)
- [Initial Setup](#)

Viewing System Faults

System Events Overview

The event log interface has been improved to provide filtering, sorting, and exporting.

Events		Dec 26, 2018 2:38:51 PM - Jan 2, 2019 2:38:51 PM					Filter: None
Timestamp	Severity	Target	Status	Description	Event		
Jan 2, 2019 2:38:05 PM	INFORMATIONAL	dbdhcp2-dbdh_66i-15451552...	Job complete	SOURCE_STOP job for "dbdhc...	alert.jobs.complete.object		
Jan 2, 2019 2:38:05 PM	INFORMATIONAL	dbdhcp2-dbdh_66i-1545155...	Job complete	SOURCE_STOP job for "dbdhc...	alert.jobs.complete.object		
Dec 30, 2018 10:37:11 AM	INFORMATIONAL	SuperHero Dev/Avengers DB ...	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 30, 2018 10:37:11 AM	INFORMATIONAL	SuperHero Dev/Shield DB 2.0	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 27, 2018 7:36:05 PM	INFORMATIONAL	SuperHero Prod/Cyborg 1.0	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 27, 2018 5:57:35 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 27, 2018 12:43:49 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 27, 2018 10:37:16 AM	INFORMATIONAL	SuperHero Dev/Shield DB 2.0	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 27, 2018 10:37:15 AM	INFORMATIONAL	SuperHero Dev/Avengers DB ...	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 26, 2018 7:05:43 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 26, 2018 3:22:03 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 26, 2018 2:40:47 PM	INFORMATIONAL	SuperHero Dev/Shield DB 2.0	Job complete	DB_REFRESH job for "SuperH...	alert.jobs.complete.object		
Dec 26, 2018 2:40:47 PM	INFORMATIONAL	SuperHero Dev/Shield DB 2.0	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 26, 2018 2:40:36 PM	INFORMATIONAL	dbdhcp2-dbdh_66i-1545155...	Job complete	ORACLE_UPDATE_REDOLOG...	alert.jobs.complete.object		
Dec 26, 2018 2:39:51 PM	INFORMATIONAL	SuperHero Dev/Avengers DB ...	Job complete	DB_REFRESH job for "SuperH...	alert.jobs.complete.object		
Dec 26, 2018 2:39:51 PM	INFORMATIONAL	SuperHero Dev/Avengers DB ...	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 26, 2018 2:39:47 PM	INFORMATIONAL	SuperHero Prod/Boomerang	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		
Dec 26, 2018 2:39:34 PM	INFORMATIONAL	SuperHero Prod/Cyborg 1.0	Job complete	DB_REFRESH job for "SuperH...	alert.jobs.complete.object		
Dec 26, 2018 2:39:34 PM	INFORMATIONAL	SuperHero Prod/Cyborg 1.0	Job complete	DB_SYNC job for "SuperHero ...	alert.jobs.complete.object		

Event Viewer screen

As shown above, the **Event Viewer** window provides information about all the events that occurred for the selected time period. Text matching is limited to the following columns:

- Action
- Description

Events		Dec 26, 2018 2:38:51 PM - Jan 2, 2019 2:38:51 PM					Filter: Warning
Timestamp	Severity	Target	Status	Description	Event		
Dec 27, 2018 5:57:35 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 27, 2018 12:43:49 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 26, 2018 7:05:43 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		
Dec 26, 2018 3:22:03 PM	WARNING	system	Unexpected management rest...	The management service is st...	alert.system.startup.manage...		

Event Viewer screen filtered for warning events

In the **Event Viewer** window, you can:

- 1 Enter filter text to reduce the results to only those rows matching the text entered. In the example above, we are filtering for "warning."
- 2 Click the **Export** button to export your results to a .csv file.
- 3 Click a column header to sort rows by the values found in that column.

The first time you click a header, rows will sort in ascending order. Clicking the same header a second time will sort the rows in descending order. Clicking the same header a third time will restore the results to their default sort order.

Procedure

1. Launch the **Delphix Management** application.
2. Click **System**.
3. Select **Events**.
4. Select a time range.

Sorting and Filtering

Optional: You can enter filter text to reduce the results to only those rows matching the text entered.

Text matching is limited to the following columns:

- Severity
- Status
- Description

You can click on a table column header to sort rows by the values found in that column.

The first time you click a header, rows will sort in ascending order. Clicking the same header a second time will sort the rows in descending order. Clicking the same header a third time will restore the results to their default sort order.

Click the page navigation buttons to advance through large result sets.

Column Resizing and Tooltips

If you wish, you can resize column widths to better fit the data to the available screen space. To resize a column:

1. Hover the mouse over a column separator found in the header. This will cause the mouse pointer to change shape.
2. Click and drag and the column separator to the desired position. Dragging to the left will reduce the column width. Dragging to the right will increase the width.
3. Release the mouse button.

The screenshot shows an 'Audit' window with a time range filter set to 'Jan 2, 2019 1:45:54 PM - Jan 2, 2019 2:45:54 PM'. The table has the following data:

Timestamp	User/Policy	Action	Description
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 1:47:13 PM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".

Alternatively, you can auto-size a column to fit the widest value of the current page:

1. Hover the mouse over a column separator found in the header. This will cause the mouse pointer to change shape.
2. Double click the column separator.

Values that do not fit within their column will be truncated with an ellipses (...). Hover the mouse over any value to see a tooltip rendering the complete, non-truncated value.

Jan 2, 2019 1:53:09 PM - Jan 2, 2019 2:53:09 PM Filter: None 

Timestamp	User/Policy	Action	Description
Jan 2, 2019 2:40:30 PM	Refresh	DB_SYNC	Run SnapSync for database "Shield DB 2.0".
Jan 2, 2019 2:40:03 PM	Refresh	ORACLE_UPDATE_REDOLOGS	Update Oracle online redo log files for virtual da...
Jan 2, 2019 2:39:36 PM	Refresh	DB_SYNC	Run SnapSyn
Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset database "dbdhcp2--dbdh_GGT-1545155338795"
Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2--dbdh_66f-15451552301..."
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Shield DB 2.0".
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Avengers DB 1.1".
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Cyborg 1.0".
Jan 2, 2019 2:27:45 PM	admin	SOURCE_DISABLE	Disable dataset "dbdhcp2-bbdhcp-AHCI-58503..."
Jan 2, 2019 2:27:44 PM	admin	SOURCES_DISABLE	Disable a list of datasets in environment "Hawk".
Jan 2, 2019 2:27:43 PM	admin	ENVIRONMENT_ENABLE	Enable environment "Hawk".
Jan 2, 2019 2:27:14 PM	admin	ENVIRONMENT_DISABLE	Disable environment "Hawk".
Jan 2, 2019 2:23:17 PM	admin	MASKINGJOB_FETCH	Fetching all Masking Jobs from "localhost".
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".

Hover over a description to see a tooltip. 

Exporting Results

Click the **Export** button to download the current page of results to a file of comma-separated values (CSV).

Related Topics

- [System Faults](#)

Accessing Audit Logs

This topic describes how to access audit logs. The audit log provides a record of all actions that were initiated by a policy or user, regardless of whether that action was successful.

Overview

Audit logs provide a record of all actions that were initiated by a policy or user, regardless of whether that action was successful. The audit log interface has been improved to provide filtering, sorting, and exporting.

Audit		Jan 2, 2019 1:53:09 PM - Jan 2, 2019 2:53:09 PM			Filter: None
Predefined Range 1 Hour Custom Range From: Jan 2, 2019 1:53:09 PM To: Jan 2, 2019 2:53:09 PM	Timestamp	User/Policy	Action	Description	
	Jan 2, 2019 2:40:30 PM	Refresh	DB_SYNC	Run SnapSync for database "Shield DB 2.0".	
	Jan 2, 2019 2:40:03 PM	Refresh	ORACLE_UPDATE_REDOLOGS	Update Oracle online redo log files for virtual da...	
	Jan 2, 2019 2:39:36 PM	Refresh	DB_SYNC	Run SnapSync for database "Avengers DB 1.1".	
	Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2-dbdh_GGT1545155338...	
	Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2-dbdh_661-15451552301...	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Shield DB 2.0".	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Avengers DB 1.1".	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Cyborg 1.0".	
	Jan 2, 2019 2:27:45 PM	admin	SOURCE_DISABLE	Disable dataset "dbdhcp2-bbdhcp-AHCI-58503...	
	Jan 2, 2019 2:27:44 PM	admin	SOURCES_DISABLE	Disable a list of datasets in environment "Hawk".	
	Jan 2, 2019 2:27:43 PM	admin	ENVIRONMENT_ENABLE	Enable environment "Hawk".	
	Jan 2, 2019 2:27:14 PM	admin	ENVIRONMENT_DISABLE	Disable environment "Hawk".	
	Jan 2, 2019 2:23:17 PM	admin	MASKINGJOB_FETCH	Fetching all Masking Jobs from "localhost".	
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".		

Audit window

As shown above, the **Audit** window displays all actions that were initiated for the selected period of time. You can enter filter text to reduce the results to only those rows matching the text entered. In the figure below, we are filtering for "user."

Text matching is limited to the following columns:

- Action
- Description

Audit		Jan 2, 2019 1:53:09 PM - Jan 2, 2019 2:53:09 PM			Filter: None
Predefined Range 1 Hour Custom Range From: Jan 2, 2019 1:53:09 PM To: Jan 2, 2019 2:53:09 PM	Timestamp	User/Policy	Action	Description	
	Jan 2, 2019 2:40:30 PM	Refresh	DB_SYNC	Run SnapSync for database "Shield DB 2.0".	
	Jan 2, 2019 2:40:03 PM	Refresh	ORACLE_UPDATE_REDOLOGS	Update Oracle online redo log files for virtual da...	
	Jan 2, 2019 2:39:36 PM	Refresh	DB_SYNC	Run SnapSync for database "Avengers DB 1.1".	
	Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2-dbdh_GGT1545155338...	
	Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2-dbdh_661-15451552301...	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Shield DB 2.0".	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Avengers DB 1.1".	
	Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Cyborg 1.0".	
	Jan 2, 2019 2:27:45 PM	admin	SOURCE_DISABLE	Disable dataset "dbdhcp2-bbdhcp-AHCI-58503...	
	Jan 2, 2019 2:27:44 PM	admin	SOURCES_DISABLE	Disable a list of datasets in environment "Hawk".	
	Jan 2, 2019 2:27:43 PM	admin	ENVIRONMENT_ENABLE	Enable environment "Hawk".	
	Jan 2, 2019 2:27:14 PM	admin	ENVIRONMENT_DISABLE	Disable environment "Hawk".	
	Jan 2, 2019 2:23:17 PM	admin	MASKINGJOB_FETCH	Fetching all Masking Jobs from "localhost".	
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".		

Audit window filtered for User events

1

Enter filter text to reduce the results to only those rows matching the text entered. In the example above, we are filtering for "user."

2

Click the **Export** icon to export your results to a .csv file.

3

You can click a column header to sort rows by the values found in that column.

The first time you click a header, rows will sort in ascending order. Clicking the same header a second time will sort the rows in descending order. Clicking the same header a third time will restore the results to their default sort order.

 Actions displayed in the Actions panel or on the Audit page are kept forever.

Procedure

1. Login to the **Delphix Management** application using **admin** credentials.
2. Click **System**.
3. Select **Audit**.
4. Select an audit log time range.

Sorting and Filtering



Timestamp	User/Policy	Action	Description
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 1:47:13 PM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".
Jan 2, 2019 1:42:44 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 1:37:03 PM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".
Jan 2, 2019 12:56:18 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 11:01:30 AM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".
Jan 2, 2019 10:52:08 AM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 10:31:20 AM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".
Jan 2, 2019 9:51:04 AM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".

1

Enter filter text (e.g. USER_LOGIN) to reduce the results to only those rows matching the text entered.

2

Click the **Export** icon to export your results to a .csv file.

3

You can click on a table column header to sort rows by the values found in that column. The first time you click a header, rows will sort in ascending order. Clicking the same header a second time will sort the rows in descending order. Clicking the same header a third time will restore the results to their default sort order. Text matching is limited to the following columns:

- Action
- Description

Click the page navigation buttons to advance through large result sets.

Column Resizing and Tooltips

If you wish, you can resize column widths to better fit the data to the available screen space. To resize a column:

1. Hover the mouse over a column separator found in the header. This will cause the mouse pointer to change shape.
2. Click and drag and the column separator to the desired position. Dragging to the left will reduce the column width. Dragging to the right will increase the width.
3. Release the mouse button.

Audit USER_LOGIN

Jan 2, 2019 1:45:54 PM - Jan 2, 2019 2:45:54 PM

Predefined Range
 1 Hour

Custom Range
 From: Jan 2, 2019 1:45:54 PM
 To: Jan 2, 2019 2:45:54 PM

Timestamp	User/Policy	Action	Description
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".
Jan 2, 2019 1:47:13 PM	sysadmin	USER_LOGIN	Log in as user "sysadmin" from IP "172.16.116.59".

To change the width of a column, hover over the column separator.

Alternatively, you can auto-size a column to fit the widest value of the current page:

1. Hover the mouse over a column separator found in the header.
This will cause the mouse pointer to change shape.
2. Double click the column separator.

Values that do not fit within their column will be truncated with ellipses (...). Hover the mouse over any value to see a tooltip rendering the complete, non-truncated value.

Jan 2, 2019 1:53:09 PM - Jan 2, 2019 2:53:09 PM Filter: None

Hover over a description to see a tooltip.

Timestamp	User/Policy	Action	Description
Jan 2, 2019 2:40:30 PM	Refresh	DB_SYNC	Run SnapSync for database "Shield DB 2.0".
Jan 2, 2019 2:40:03 PM	Refresh	ORACLE_UPDATE_REDOLOGS	Update Oracle online redo log files for virtual da...
Jan 2, 2019 2:39:36 PM	Refresh	DB_SYNC	Run SnapSyn...
Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2-- database "dbdhcp2-- dbdh_GGT-1545155338795".
Jan 2, 2019 2:38:01 PM	Refresh	SOURCE_STOP	Stop dataset "dbdhcp2--dbdh_66f1-15451552301...
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Shield DB 2.0".
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Avengers DB 1.1".
Jan 2, 2019 2:38:00 PM	Refresh	DB_REFRESH	Refresh database "Cyborg 1.0".
Jan 2, 2019 2:27:45 PM	admin	SOURCE_DISABLE	Disable dataset "dbdhcp2-bbdhcp-AHCI-58503...
Jan 2, 2019 2:27:44 PM	admin	SOURCES_DISABLE	Disable a list of datasets in environment "Hawk".
Jan 2, 2019 2:27:43 PM	admin	ENVIRONMENT_ENABLE	Enable environment "Hawk".
Jan 2, 2019 2:27:14 PM	admin	ENVIRONMENT_DISABLE	Disable environment "Hawk".
Jan 2, 2019 2:23:17 PM	admin	MASKINGJOB_FETCH	Fetching all Masking Jobs from "localhost".
Jan 2, 2019 2:23:14 PM	admin	USER_LOGIN	Log in as user "admin" from IP "172.16.116.59".

Exporting Results

Click the icon to download the current page of results to a file of comma-separated values (CSV).

Related Topics

- [Monitoring and Log Management](#)
- [Creating Support Logs](#)
- [Setting Support Access Control](#)

Creating Support Logs

This topic describes how to create support bundles and manage server access control for Delphix Support.

Support bundles are used by Delphix as diagnostic tools for resolving Delphix Engine issues and routine maintenance. Support bundles can be transferred directly to Delphix Support or downloaded. No customer-specific data is included in the support bundle information. All passwords and personal data are either encrypted or omitted. This is an outbound-only connection from the Delphix Engine.

Using the GUI

1. Log into the **Delphix Management** appliance as an Engine administrator.
2. Click **Help**.
3. Select **Support Logs**.
4. The **Support Bundle** dialog appears.

Support Bundle [X]

Select Support Bundle Type

All

Download or Transfer Support Bundle

Transfer

Download

Support Case Number

Hide advanced ^

Analytics

Include all analytics data (the default is a 10MB limit)

Close OK

5. Select **Download** or **Transfer**.
 - a. If you select **Download**, then the support logs will be downloaded as a compressed ".tar" file into a folder on your workstation.
 - b. If you select **Transfer**, then the support logs will be uploaded over HTTPS to Delphix Support. If you have configured an HTTP Proxy, it will be used to send the support logs.

- c. If there is a support case involved, then please enter the case number to associate the logs to the case.
6. Click on **Show advanced** and select **Analytics**. This will include all the analytics data (default, up to 10MB) in the Support Bundle.
7. Click **OK**.
 - a. If you selected **Download** and have the compressed ".tar" file in a folder on your workstation, please **upload** that file to Delphix Support via the website at <http://upload.delphix.com>.
 - b. If there is a support case involved, then please enter the case number (again) to associate the logs to the case.

You can also access support log functionality in the **ServerSetup** application using **sysadmin** credentials. Click **Support Bundle** in the top menu bar.

Using the CLI

1. ssh into your Delphix Engine.

```
ssh <sysadmin_user>@<delphixengine>
```

2. Run the upload operation.

```
delphix > service
delphix service > support
delphix service support > bundle
delphix service support bundle > upload
```

3. Commit the operation.

```
delphix service support bundle upload *> commit
```

Related Topics

- [Accessing Audit Logs](#)
- [Setting Support Access Control](#)

Setting Support Access Control

This topic describes how to set the Support Access Control for Delphix Support. Support access control enables Delphix Support to access your instance of the Delphix Engine for a defined period of time using an access token.

1. Log into the **Delphix Setup** application using **sysadmin** credentials.
2. Click **Server Preferences**.
3. Select **Support Access**.
4. Click **Enable**.
5. Set the time period during which you want to allow Delphix Support to have access to your instance of the Delphix Engine.
6. Click **Generate Token**.
Provide the token to Delphix Support to enable access to your server.

Related Topics

- [Creating Support Logs](#)

Setting Syslog Preferences

Syslog is a widely used standard for message logging. It permits the separation of the software that generates messages, the system that stores them, and the software that reports and analyzes them. Delphix makes use of Syslog as one of the standard mechanisms, along with SNMP and email, to distribute important user and system events, such as alerts, faults, and audits. In the case of Delphix, each Delphix Engine acts as a Syslog client which propagates the events to a centralized Syslog server.


The network protocol over which the Delphix Engine communicates with the Syslog server is standardized in [RFC 5424](#). As a protocol, it supports using either UDP ([RFC 5426](#)) or TCP ([RFC 6587](#)) as the underlying transport and optional TLS mapping has been introduced to encrypt the messages over the wire for security purposes ([RFC 5425](#)). However, as of this release, we only support Syslog over UDP with no encryption, which implies that Syslog messages are always sent in the clear and may be lost during transmission and delivered out of order due to the limitations of UDP.

To configure for Syslog support, you must specify the communication endpoint to which the Syslog server listens, which includes the hostname or IP address of the Syslog server and an optional port number. The latter defaults to 514 according to the Syslog standard but it can be changed if necessary.

System and user events generated by Delphix are always forwarded immediately to the Syslog server, which ensures the timely delivery of important events that may require immediate action.


A couple of different output formats are supported for messages delivered over Syslog, namely, TEXT and JSON. The TEXT format is the default. To change the message format, as of this release, you must do so via the CLI.

1. Log into the **Delphix Setup** application using **sysadmin** credentials.
2. Select **Preferences > Syslog Configuration**.



3. Select the severity level of the messages you want to be sent to the Syslog server.
4. Click the **pencil** icon next to **Syslog Servers** and then in the **Syslog Configuration** window select the plus icon.

5. Enter the Syslog server hostname/IP address and port number.



The image shows a 'Syslog Configuration' dialog box. At the top, the title 'Syslog Configuration' is displayed with a close button (X) on the right. Below the title, there is a 'Syslog Severity' section with a dropdown menu currently set to 'Warning'. The main part of the dialog is a 'Syslog Servers' section, which contains a sub-dialog for adding a server. This sub-dialog has two input fields: 'Address' with the value '10.10.10.1' and 'Port' with the value '514'. The sub-dialog has a close button (X) and a confirmation button (blue with a checkmark). Below the sub-dialog, there is a 'Syslog Status' section with a checkbox labeled 'Enable Syslog', which is currently unchecked. At the bottom right of the main dialog, there are two buttons: 'Cancel' and 'Save'.

6. Select **Enable Syslog**.
7. Click **Save**.

Related Topics

- [Monitoring and Log Management](#)
- [Severity Levels for Syslog Messages](#)

Severity Levels for Syslog Messages

This topic discusses the Syslog reporting feature of the Delphix Engine, along with severity levels.

Syslog is a widely used standard for message logging. It permits the separation of the software that generates messages, the system that stores them, and the software that reports and analyzes them. Delphix makes use of Syslog as one of the standard mechanisms, along with SNMP and email, to distribute important user and system events, such as alerts, faults, and audits. In the case of Delphix, each Delphix Engine acts as a Syslog client which propagates the events to a centralized Syslog server.

Every Syslog message is attached to a severity level. As the name suggests, the severity level describes the severity of the event in question.

Audit Logs

Audit records are Informational Syslog messages. If you would like to forward Audit records, choose Severity Level Informational.

Severity levels

Every Syslog message is attached to a severity level number. Delphix defines the severity of Syslog messages in accordance with RFC 3164. There are eight severity levels available, as follows:

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

When setting up the Syslog settings for your Delphix Engine, you have the ability to choose what alerts to report. The severity levels above are available for users to select. Once you select a severity level, the Delphix Engine will send messages of the same or higher severity (i.e., the same or lower number) to your Syslog server. Therefore, there is no reason to select more than one severity. For example, if the "Notice" severity level is selected, all events less severe than Notice (Informational and Debug) will not be reported. If you want all events to be reported via Syslog, the Debug severity level should be chosen.

Related Topics

- [Setting Syslog Preferences](#)
- [Monitoring and Log Management](#)


Diagnosing Connectivity Errors

Prior to the 5.1 release, when the Delphix Engine ran into an error operating on an external database or environment, it reported the immediate error that it had encountered; there was no mechanism for automatic analysis of the root causes of failures. The 5.1 release included infrastructure for automatic diagnosis of errors. When one of these errors occurs, the Delphix Engine now launches a set of tests to locate the root cause of the problem and present the result of the diagnosis. This will help you easily identify the true sources of errors such as closed ports or misconfigured router.


Failed Actions

The Delphix Engine communicates failures in two different manners: actions that fail to complete, and faults. To view failed actions:

1. In the top right-hand corner of the Delphix Management application, click **Actions**.
2. For more information about why the action failed, click the (i) icon to show the error dialog.



The following shows a popup message with more information about the problem and what actions to take to resolve it. For some errors, the Delphix Engine will be able to diagnose the problem further and display this extra information under **Diagnosing Information**. In the screenshot above, the job failed because the Delphix Engine was unable to lookup the host address.



Viewing Active Faults

A fault symbolizes a condition that can affect the performance or functionality of the Delphix Engine and must be addressed. Faults can be either warnings or critical failures that prevent the Delphix Engine from functioning normally. For example, a problem with a source or target environment can cause SnapSync or LogSync policy jobs to fail. Faults will show up as active as long as:

- The error is still occurring, or
- You have chosen to manually resolve it or ignore it

For example, if a background job fails, it will create a fault that describes the problem. To view any active faults:

- In the top right-hand corner of the Delphix Management interface, click **Faults**.

This brings the Faults screen listing all active faults.

Faults

Current Archive

Refresh (Manual) Resolve All Filter: none ↑

Severity	Diagnosed	Title	Target
<input checked="" type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:43...	Incorrect toolkit owner	hana-virtual-tgt.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:43...	Incorrect toolkit owner	hana-virtual-src.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:38...	Incorrect toolkit owner	postgres-virtual-tgt.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:38...	Incorrect toolkit owner	postgres-virtual-src.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:35...	Incorrect toolkit owner	db2-virtual-tgt.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 29, 2021 12:34...	Incorrect toolkit owner	db2-virtual-src.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 9:38 ...	Failed to start database	EmptyVDB_CUFMT6SQ
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 9:36 ...	Invalid database credentials	DBOMSRBBD6C
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 9:00 ...	Failed to connect to the source database	CDOMLOTG4F5E:UNKNOWN:6f1hSg
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 9:00 ...	Incorrect toolkit owner	mg-centos-75-oracle-18000-tgt.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 8:59 ...	Failed to connect to the source database	CDOMLOS RB5A3:UNKNOWN:zWmLFU
<input type="checkbox"/> ▲ WARNING	Sep 28, 2021 8:59 ...	Incorrect toolkit owner	mg-centos-75-oracle-18000-src.dlpxdc.co
<input type="checkbox"/> ▲ WARNING	Sep 27, 2021 5:07 ...	Unable to purge logs	Untitled/CDOMLOTG4F5E@2021-09-20T07:56:26.571Z

▲ WARNING

Date
Sep 29, 2021 12:43 PM

Title
Incorrect toolkit owner

Target
hana-virtual-tgt.dlpxdc.co

Details
The owner of the Delphix toolkit installation directory '/var/tmp' is not 'hdbadm'.

User Action
Check the owner of the toolkit directory.

● Resolving and Ignoring Resolve Ignore

1 to 21 of 21 << < Page 1 of 1 > >>

The screenshot above illustrates a fault with regard to a failure to a TCP slot table entry. The Delphix Engine will mark an object with a warning triangle to indicate that it is affected by an external problem. You can view more details of the fault by looking at the active faults and their fault effects.

Related Topics

- [System Faults](#)

Email (SMTP) Alert Notifications

Overview


The configuration for SMTP-based alert notifications has two components:

- The configuration of an SMTP gateway by the Delphix system administrator
- The configuration of one or more alert profiles (if needed)

Configuring the SMTP Gateway

Before email-based alerts can function properly, many organizations require that an SMTP gateway is configured, through which all outbound email is sent.

1. Contact the appropriate administrator for your site in order to determine the SMTP gateway settings.
2. Login to the **Delphix Setup** application as **sysadmin** or another user with system administrator privileges.
3. On the Dashboard screen, locate the **Outbound Connectivity** section, and click **Modify**.



Outbound Connectivity for SMTP Gateway

4. If not checked already, check the box next to **Use an existing SMTP server**.

Outbound Connectivity

WEB PROXY
The Web Proxy Server will be used to communicate with Delphix Corp. for support, troubleshooting, upgrades, updates, and patches.

Configure web proxy

PHONE HOME SERVICE
If enabled, this service will automatically send a minimal support bundle once a day to the Delphix support site over HTTPS. This will help with future support and troubleshooting. A connection to the internet, either directly or via web proxy is required.

Enable phone home service

USER-CLICK ANALYTICS
If enabled, this service will automatically send a stream of anonymous, non-personal metadata describing user interaction with the product's user interface. This data will help us to better understand how our products are being used, and to improve our products and services.

Enable Usage Analytics

SMTP
Configure the Delphix Engine's SMTP sending service to enable email notifications. Your sysadmin email will be used for receiving system reports, events, and fault notifications.

Use an existing SMTP server
Enable email notifications for faults and events

Server Name or IP Address

Port

25

TLS Authentication

SMTP Authentication

From Email Address

noreply@delphix.com

SMTP Send Timeout

60

Maximum timeout to wait, in seconds, when sending mail.

Test Email Address

Comma-separated list of Test Email Address

Cancel
Save

5. At a minimum, enter the required information:
- a. SMTP Server Name of IP Address
 - b. SMTP port
 - c. From Email Address
This will be the email address from which all alert emails will be sent.
6. In **Test Email Address**, enter the same email address to verify that you are able to receive email properly.
7. Click **Save** to save changes.

For further information, see the “Outbound Connectivity” section of [Initial Setup](#).

Alert Profiles

The Delphix Engine can send out email notifications when alerts happen. Alert profiles control this functionality.

An alert profile is composed of two things:

- **Filter Specification:** A filter, or combination of filters, that specifies which alerts are of interest.
- **Alert Action:** This specifies the email addresses to which the Delphix Engine will send an email when an alert matches the filter specification.

By default, the Delphix Engine has a single alert profile configured with the following parameters:

- Filter Specification: Match any alert with a severity level of **CRITICAL** or **WARNING**.
- Alert Actions: Send an email to the address defined for user **admin**.

i Default domain user

The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Using the CLI, it is possible to:

- Modify the system default alert profile
- Create additional profiles in addition to the default one
- Set multiple actions for a single profile, such as "email delphix_admin" and "email user1@mycompany.com".

Simple Filters

- Filtered by Owner of alerts target – for example, objects owned by user 1

Complex Filters

Complex filters combine/modify other sub-filters:

- “And” filter – Used when all conditions defined must be met for the filter to notify the user with an email
- “Or” filter – Used when either one or the other of the conditions defined in the filters must be met for the filter to notify the user with an email
- “Not” filter – Used to exclude items

Limitations

- This is a CLI feature.
- Alert Profiles do not override permission settings. If you do not have Read permission on an object then your alert profile will never get triggered for that object's alerts, regardless of your filter settings.

The following CLI examples will run through how to create these three filters. Each example provides three different methods of setting up a profile. These include the following:

- A simple profile
- A profile with two filters
- A complicated profile

For more information, see [CLI Cookbook: Creating Alert Profiles](#).

A Simple Profile

A simple profile approach matches the Delphix out-of-the-box default alert profiles. To create a simple alert profile using the CLI as seen in the figure below, go into the alert profile section of the command-line interface (CLI) and create a new profile. Line four prompts the engine to send an email when the filters are triggered. The following three command lines refer to the filter specifications. Follow two severity levels: warning and critical. This will trigger an email alert when any warning or critical events occur.

```
delphix > cd alert
delphix alert > cd profile
delphix alert profile > create
delphix alert profile create > set actions.0.type=AlertActionEmailUser
```

```
delphix alert profile create > set filterSpec.type=SeverityFilter
delphix alert profile create > set filterSpec.severityLevels.0=CRITICAL
delphix alert profile create > set filterSpec.severityLevels.1=WARNING
delphix alert profile create > commit
```

Simple Alert Profile example in the CLI

A Compound Alert Profile

Creating a compound alert profile in the CLI will combine two filters together. This profile triggers an email about any alert on objects owned by the delphix_admin plus any other alert that is critical. The compound alert profile creates two filters. The first one will be the target owner filter, which in this case is **admin**. The second filter is the severity filter, allowing users to match anything that is critical. Combine these two filters using the OR logic so that if any of the sub-filters match, the whole filter matches. An example of this can be seen in the figure below.



Alert Profile using OR logic

While working in the CLI, the first four lines describe a simple alert profile. The distinction between simple and compound alert profiles is that in a compound profile, the top-level filter uses an OR filter with sub-filters for target owner and severity level, as seen in line five of the figure below.

```
delphix > cd alert
delphix alert > cd profile
delphix alert profile > create
delphix alert profile create > set actions.0.type=AlertActionEmailUser
delphix alert profile create > set filterSpec.type=OrFilter
delphix alert profile create > set filterSpec.subFilters.0.type=TargetOwnerFilter
delphix alert profile create > set filterSpec.subFilters.0.owners.0=delphix_admin
delphix alert profile create > set filterSpec.subFilters.1.type=SeverityFilter
delphix alert profile create > set filterSpec.subFilters.1.severityLevels.0=CRITICAL
delphix alert profile create > commit
```

A Compound Alert Profile

Complex Alert Profile

A complex alert profile uses the profile filter created in the compound alert profile and modifies it. For the example shown in the figure below, you have a VDB named `test_instance` that you do not need any emails about. The following commands will create an effective filter.

1. Create an OR filter with two sub filters: target owner and event type.
2. Create a NOT filter that will exclude the VDB (`test_instance`) from which you do not want to receive notifications.
3. Use the AND logic to combine all these filters together, as seen below.



Complex Alert Profile

Below is an example of the command lines used to set up this complex profile.

```

delphix > cd alert
delphix alert > cd profile
delphix alert profile > create
delphix alert profile create > set actions.0.type=AlertActionEmailUser
delphix alert profile create > set filterSpec.type=AndFilter
delphix alert profile create > set filterSpec.subFilters.0.type=NotFilter
delphix alert profile create > edit filterSpec.subFilters.0.subFilter
delphix alert profile create filterSpec.subFilters.0.subFilter > set
type=TargetFilter
delphix alert profile create filterSpec.subFilters.0.subFilter > set targets.0=test_i
nstance
delphix alert profile create filterSpec.subFilters.0.subFilter > back
delphix alert profile create > set filterSpec.subFilters.1.type=OrFilter
delphix alert profile create > set filterSpec.subFilters.1.subFilters.0.type=TargetOw
nerFilter
delphix alert profile create > set filterSpec.subFilters.1.subFilters.0.owners.0=delp
hix_admin
  
```

```
delphix alert profile create > set filterSpec.subFilters.1.subFilters.1.type=Severity
Filter
delphix alert profile create > set filterSpec.subFilters.1.subFilters.1.severityLevel
s=CRITICAL
delphix alert profile create > commit
```

Complex Alert Profile CLI

Creating Alert Profiles

1. SSH into your engine's CLI using your `delphix_admin` username and password

```
ssh delphix_admin@yourdelphixengine
```

2. Start creating your new profile. After logging in to get to the alert section enter `alert`.

```
delphix > alert
delphix alert > profile
delphix alert profile > create
delphix alert profile create > ls
```

3. Set Action(s)

Use **AlertActionEmailList** if you want to specify a list of email addresses for this profile.

```
delphix alert profile create > set actions.0.type=AlertActionEmailList
delphix alert profile create > set actions.0.addresses.0=<email address to send
to>
delphix alert profile create > set actions.0.addresses.1=<additional email
address>
delphix alert profile create > set actions.0.addresses.2=<additional email
address>
```

Or, use **AlertActionEmailUser** if you just want the emails to go to the email address associated with this Delphix user.

```
delphix alert profile create > set actions.0.type=AlertActionEmailUser
```

It is possible to add more than one action here, so you may use both **AlertActionEmailList** and **AlertActionEmailUser** if desired.

4. Set filter

Here is an example of setting a simple severity filter. With this filter, emails will be sent for any **CRITICAL** or **WARNING** alerts.

```
delphix alert profile create > set filterSpec.type=SeverityFilter
delphix alert profile create > set filterSpec.severityLevels.0=CRITICAL
delphix alert profile create > set filterSpec.severityLevels.1=WARNING
```

Here is an example of setting a simple target-owner filter. With this filter, emails will be sent for any alert whose target is owned by `delphix_admin`.

```
delphix alert profile create > set filterSpec.type=TargetOwnerFilter
delphix alert profile create > set filterSpec.owners.0=delphix_admin
```

Here is an example of a compound filter. With this filter, we combine the above two filters – an email is sent when an alert is **CRITICAL** or **WARNING**, and the alert's target is owned by delphix_admin.

```
delphix alert profile create > set filterSpec.type=AndFilter
delphix alert profile create > set filterSpec.subFilters.0.type=SeverityFilter
delphix alert profile create > set filterSpec.subFilters.0.severityLevels.0=CRITICAL
delphix alert profile create > set filterSpec.subFilters.0.severityLevels.1=WARNING
delphix alert profile create > set filterSpec.subFilters.1.type=TargetOwnerFilter
delphix alert profile create > set filterSpec.subFilters.1.owners.0=delphix_admin
```

5. Commit your changes

```
delphix alert profile create > commit
```

Profile Filters

As seen above, you can use different filter types to customize which alerts the Delphix Engine will send emails about. The various filter types are listed below.

Simple Filters

Filter Type	Purpose	Example	Allowed Values
SeverityFilter	Match based on the alert's severity level (critical, warning, informational)	severityLevels.0=CRITICAL severityLevels.1=WARNING This would match any alert whose severity level is CRITICAL or WARNING.	1 or more of: <ul style="list-style-type: none"> CRITICAL WARNING INFORMATIONAL
EventFilter	Match based on the alert's event type.	eventTypes.0=fault.* This would match any alert that is generated due to a newly-raised fault on the engine.	One or more text entries, optionally using the * wildcard.

Filter Type	Purpose	Example	Allowed Values
TargetFilter	Match based on the alert's target.	targets.0="Group/DB" This would match any alert whose target is the database "DB" located in the group "Group".	Any object in the system. 1 or more objects may be specified.
TargetOwnerFilter	Match based on the owner of the alert's target.	owners.0=delphix_admin This would match any alert whose target's owner is the delphix_admin user.	Any user in the system. 1 or more users may be specified.

Compound Filters

These filters combine/modify the behavior of other filters, called "subfilters". The subfilters may be of any type (simple or complex).

Filter Type	Purpose	Number of subfilters required
AndFilter	This filter matches if all subfilters match	2 or more
OrFilters	This filter matches if any subfilter matches.	2 or more
NotFilter	This filter matches if the subfilter does not match.	1

Action Types

With the AlertActionEmailUser type, notification emails will be sent to the email address of the user who owns the alert profile.

With the AlertActionEmailList type, a list of email addresses must be specified in the "addresses" array. Notification emails will be sent to these addresses.

Email Format Options

You can send plain text as well as structured JSON. JSON can be useful for constructing solutions that will parse the email and perform further actions (notify, escalate, log).

To change the format, while updating an alert profile:

```
delphix alert profile ALERT_PROFILE-X update > set actions.0.format=<JSON|TEXT>
```

Related Topics

- [CLI Cookbook: Creating Alert Profiles](#)

- [Monitoring and Log Management](#)

Splunk Integration

Delphix enables self-monitoring/diagnosability of Delphix Engines by providing native integration with Splunk Enterprise. By providing details about your Splunk instance, you can allow Delphix to send structured JSON logs to Splunk that capture activity on the Delphix engine(s). These logs include Delphix events (Actions, Job Events, Faults, and Alerts) as well as performance metrics (CPU, disk, network, TCP, dataset, NFS, iSCSI) and capacity metrics. Delphix Insight enables extensible search and visualization of actionable information and provides a centralized, comprehensive view of Delphix activity (including the ability to cross-reference information from multiple Delphix engines) on a platform that allows building your own operational intelligence for your Delphix installation.

This section covers the following topics:

- [Configuring Splunk](#)
- [Using Search](#)

Configuring Splunk

Prerequisites

Before you configure the Delphix Engine you will need to configure and make a note of the following in Splunk:

 Please refer to the [Splunk documentation](#) for detailed steps on how to configure your values.

Supported Splunk versions

Delphix only supports Splunk Enterprise 6.3.0 or higher.

1. In the Splunk web **UI Enable SSL** (this is optional but best practice for security) in your global HTTP Event Collector (HEC) settings.
2. The **Splunk hostname** or **IP Address**.
3. The **HEC Port** number for your Splunk instance (default 8088).
4. Enable the HTTP Event Collector on Splunk, and create a new **HEC Token** with a new Splunk index set as an allowed index for the token. Make sure **Enable Indexer Acknowledgement** is **unchecked** for the token.

 If you wish, you can use a separate Splunk index for performance and capacity metrics (otherwise, the same index will be used for both events and metrics). If you are using Splunk 7.0+, it is recommended that you create this second index as a special “Metrics” type index that is optimized for indexing and searching metrics data.


5. Note the **HEC Token Value** and the **Allowed Indexes** for the token.

The following table provides an example of the data you will need.

Attribute	Sample Data
Splunk Server IP address	192.168.8.8
Splunk Server HEC Port Number	8088
Splunk HEC Token	12345678-1234-1234-1234-1234567890AB
Index Name for Events	delphix_events
Index Name for Metrics	delphix_metrics

Configuring Delphix for Splunk

1. Log in to the **Delphix Server Setup UI** as the sysadmin.
2. From the **Preferences** menu select **Splunk Configuration**.



3. In the Splunk Configuration window, enter your Splunk values. To reduce the volume of data that will be sent to Splunk, you can optionally uncheck **Enable Metrics**.

A screenshot of the 'Splunk Configuration' window. The window has a title bar with 'Splunk Configuration' and a close button (X). The content is organized into sections: 'HTTP Event Collector Configuration' with a checkbox for 'Enable Splunk Configuration' (unchecked), fields for 'Host', 'HEC Port', 'HEC Token', 'Main Index', and 'Events Push Frequency (seconds)' (set to 60), and a checkbox for 'Enable SSL (Must match HEC settings in Splunk)'. The 'Metrics' section has a checkbox for 'Enable Metrics' (checked), a field for 'Metrics Index', a field for 'Metrics Push Frequency (seconds)' (set to 60), and a dropdown for 'Performance Data Granularity' (set to 'Per Minute'). The 'Test Connection' section has a 'Send Test Data' button. At the bottom right are 'Cancel' and 'Save' buttons.

Host	Splunk hostname or IP address
HEC Port	The TCP port number for the Splunk HTTP Event Collector (HEC)
HEC Token	The token for the Splunk HTTP Event Collector (HEC)
Main Index	The Splunk Index events will be sent to. It must be set as an allowed index for the HEC token.
Events Push Frequency	The frequency at which the Events will be pushed to Splunk. Specified in seconds.
Enable SSL	Whether to use HTTPS to connect to Splunk. Must match your HTTP Event Collector settings in Splunk.
Metrics Index	The Splunk Index metrics will be sent to. If none is specified then the Main Index will be used for metrics as well. It must be set as an allowed index for the HEC token.
Metrics Push Frequency	The frequency at which the Performance Metrics will be pushed to Splunk. Specified in seconds
Performance Data Granularity	The resolution of performance metrics data sent to Splunk. This controls the granularity of system performance metric values sent to Splunk. For example: Setting this to "Per Minute" results in a system performance metric to have a per-minute aggregated value (as collected by the analytics subsystem of the Delphix Engine) instead of a per-second value.

 **Reducing the volume of data**

To reduce the volume of data that will be sent to Splunk, you can change the Performance Data Granularity to Per Minute or uncheck Enable Metrics.

4. Click **Send Test Data** to verify your provided values. This will send a test event to the provided token and indexes.
5. Click **Save** to enable the Splunk configuration and begin sending all new Actions, Job Events, Faults, Alerts, and Metrics to your Splunk instance.

Related Topics

- [Splunk Integration](#)
- [Using Search](#)

Using Search

Use the search to analyze your data and enumerate items in a metrics index. For more about searching a metrics index, refer to the [Splunk](#) documentation.

Search Examples - Metrics

The following examples provide information on viewing Metrics on Splunk 7.x

To get a list of all Metrics:

```
| mcatalog values(metric_name)
```

To get a list of all dimensions of a given metric - say CPU utilization percentage:

```
| mcatalog values(_dims) where metric_name="system.cpu.util.pct"
```


To view the average values of overall CPU utilization percentage across all hosts with a span of 30 seconds:

```
| mstats avg(_value) WHERE index=delphix_metrics AND metric_name=system.cpu.util.pct span=30s
```

You can also display results in a chart with CPU wildcard:

```
| mstats perc85(_value) AS val85 avg(_value) AS val where metric_name="system.cpu.*" span=1s by data.kernel, data.user, data.idle
| eval total='data.kernel' + 'data.user' + 'data.idle'
| eval sys_pct= (('data.kernel'/total) * 100)
| eval usr_pct= (('data.user'/total) * 100)
| eval idle_pct= (('data.idle'/total) * 100)
| timechart span=10m avg(val) as "cpu.overall", avg(val85) as "cpu.overall 85th Percentile", avg(sys_pct) as "cpu.system", avg(usr_pct) as "cpu.user", avg(idle_pct) as "cpu.idle"
```


This type of search can be used to stack different CPU metrics that add up to 100%. Here is a sample screenshot of the above “stack different CPU metrics” from the Delphix Engines.



Search Examples - Events

The following queries demonstrate some basic visualizations of various Delphix events. The **delphix_index** value should be replaced with the name of the **Main Index** provided in during Delphix Setup. These examples serve as useful starting points that can be expanded to include other relevant data. See [Events Format](#) for a full description of the structure of each type of event.

Display event statistics per host



```

Search

index="delphix_index" | stats
count(eval(source="delphix.events.action.completed" OR
source="delphix.events.action.started" OR
    
```

```
source="delphix.events.action.waiting" )) AS ACTIONS
count(eval(source="delphix.events.fault.posted")) AS FAULTS
count(eval(source="delphix.events.alert" )) AS ALERTS
count(eval(source="delphix.events.job.event" )) AS JOBS BY host
```

List Actions in descending order by the duration

Job Run Times		
Jobs by Duration		
reference ↕	title ↕	duration ↕
ACTION-235	DB_EXPORT	231
ACTION-42	ENVIRONMENT_DISCOVER	211
ACTION-68	ENVIRONMENT_CREATE_AND_DISCOVER	206
ACTION-70	ENVIRONMENT_CREATE	176
ACTION-160	DB_SYNC	173
ACTION-72	HOST_ADD	165
ACTION-167	DB_REFRESH	161
ACTION-162	DB_PROVISION	132
ACTION-86	ENVIRONMENT_REFRESH_AND_DISCOVER	118
ACTION-146	DB_REFRESH	108

[« prev](#)
1
2
3
4
5
6
7
8
9
10
[next »](#)

Search

```
index="delphix_index" source="delphix.events.action.*" | transaction reference |
table reference title duration | sort duration
```

Lists Faults

Faults			
host ↕	dateDiagnosed ↕	details ↕	reference ↕
ip-10-110-193-142	2018-07-05T23:34:37.968Z	The TCP sunrpc.tcp_slot_table_entries property is currently set to '16' which is below the recommended minimum value of 128.	FAULT-35
ip-10-110-193-142	2018-07-05T23:34:37.902Z	The default send buffer that can be allocated for a TCP socket dictated by the second value in the net.ipv4.tcp_wmem property is currently set to '16384' which is below the recommended default of 4194304.	FAULT-34
ip-10-110-193-142	2018-07-05T23:34:37.612Z	The default receive buffer that can be allocated for a TCP socket dictated by the second value in the net.ipv4.tcp_rmem property is currently set to '87380' which is below the recommended default of 16777216.	FAULT-33
ip-10-110-193-142	2018-07-05T23:20:45.107Z	The owner of the Delphix toolkit installation directory is not 'sybase'.	FAULT-32
ip-10-110-193-142	2018-07-05T23:19:19.827Z	The owner of the Delphix toolkit installation directory is not 'sybase'.	FAULT-31

« prev 1 2 3 4 5 6 7 next »

Search

```
index="delphix_index" source="delphix.events.fault.posted" | table host
dateDiagnosed details reference
```

Completed Jobs

Events with "Job Complete" Status			
host ↕	timestamp ↕	messageDetails ↕	parentAction ↕
ip-10-110-193-142	2018-07-05T23:36:50.018Z	SUPPORT_BUNDLE_DOWNLOAD job for "delphix_admin" completed successfully.	ACTION-1585
ip-10-110-193-142	2018-07-05T23:27:43.114Z	DB_UNDO job for "Untitled/dbdhcp3--dbdh_WSH-1530832953092" completed successfully.	ACTION-1386
ip-10-110-193-142	2018-07-05T23:27:14.739Z	SOURCE_STOP job for "dbdhcp3--dbdh_WSH-1530832953092" completed successfully.	ACTION-1387
ip-10-110-193-142	2018-07-05T23:27:02.994Z	DB_REFRESH job for "Untitled/dbdhcp3--dbdh_WSH-1530832953092" completed successfully.	ACTION-1383
ip-10-110-193-142	2018-07-05T23:27:02.977Z	DB_SYNC job for "Untitled/dbdhcp3--dbdh_WSH-1530832953092" completed successfully.	ACTION-1385

[« prev](#)
[1](#)
[2](#)
[3](#)
[4](#)
[5](#)
[6](#)
[7](#)
[8](#)
[9](#)
[10](#)
[next »](#)

Search

```
index="delphix_index" source="delphix.events.job.event" | spath jobState | search jobState=COMPLETED | table host timestamp messageDetails parentAction
```

Lists Faults

Faults			
host ↕	dateDiagnosed ↕	details ↕	reference ↕
ip-10-110-193-142	2018-07-05T23:34:37.968Z	The TCP sunrpc.tcp_slot_table_entries property is currently set to '16' which is below the recommended minimum value of 128.	FAULT-35
ip-10-110-193-142	2018-07-05T23:34:37.902Z	The default send buffer that can be allocated for a TCP socket dictated by the second value in the net.ipv4.tcp_wmem property is currently set to '16384' which is below the recommended default of 4194304.	FAULT-34
ip-10-110-193-142	2018-07-05T23:34:37.612Z	The default receive buffer that can be allocated for a TCP socket dictated by the second value in the net.ipv4.tcp_rmem property is currently set to '87380' which is below the recommended default of 16777216.	FAULT-33
ip-10-110-193-142	2018-07-05T23:20:45.107Z	The owner of the Delphix toolkit installation directory is not 'sybase'.	FAULT-32
ip-10-110-193-142	2018-07-05T23:19:19.827Z	The owner of the Delphix toolkit installation directory is not 'sybase'.	FAULT-31

« prev 1 2 3 4 5 6 7 next »

Search

```
index="delphix_index" source="delphix.events.fault.posted" | table host
dateDiagnosed details reference
```

Completed Jobs

Events with "Job Complete" Status			
host ↕	timestamp ↕	messageDetails ↕	parentAction ↕
ip-10-110-193-142	2018-07-05T23:36:50.018Z	SUPPORT_BUNDLE_DOWNLOAD job for "delphix_admin" completed successfully.	ACTION-1585
ip-10-110-193-142	2018-07-05T23:27:43.114Z	DB_UNDO job for "Untitled/dbdhcp3–dbdh_WSH-1530832953092" completed successfully.	ACTION-1386
ip-10-110-193-142	2018-07-05T23:27:14.739Z	SOURCE_STOP job for "dbdhcp3–dbdh_WSH-1530832953092" completed successfully.	ACTION-1387
ip-10-110-193-142	2018-07-05T23:27:02.994Z	DB_REFRESH job for "Untitled/dbdhcp3–dbdh_WSH-1530832953092" completed successfully.	ACTION-1383
ip-10-110-193-142	2018-07-05T23:27:02.977Z	DB_SYNC job for "Untitled/dbdhcp3–dbdh_WSH-1530832953092" completed successfully.	ACTION-1385

« prev 1 2 3 4 5 6 7 8 9 10 next »

Search

```
index="delphix_index" source="delphix.events.job.event" | spath jobState | search jobState=COMPLETED | table host timestamp messageDetails parentAction
```

Search Examples - Event Formats

The Actions, Job Events, Faults, and Alerts that Delphix sends to Splunk are structured according to predefined [JSON schemas](#).

JSON Schemas for Events

The following set of [JSON schemas](#) define the shape of each Splunk event, including which properties are expected to exist for each event type and what those properties mean. Some of these .json files are used as shared building blocks to define the other schemas; the top-level schemas which define each distinct event type are **Action.json**, **Alert.json**, **JobEvent.json**, **Fault.json**, and **FaultEffect.json**.

Below key-values are Splunk event metadata - This follows GeneralSplunkHeader.json - schema and GeneralSplunkEvent.json

GeneralSplunkHeader.json

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "time": {
      "description": "The time the event was logged. The default time format is epoch time format, in the format <sec>.<ms>.",

```

```

        "type": "number"
    },
    "host": {
        "description": "The system's hostname. Will be the host value assigned to
the event data in Splunk.",
        "type": "string"
    },
    "source": {
        "description": "For example mgmt.event.action. Will be the source value
to assigned to the event data in Splunk.",
        "type": "string"
    },
    "sourcetype": {
        "description": "The sourcetype value to assign to the event data.",
        "type": "string"
    },
    "index": {
        "description": "The name of the index by which the event data is to be
indexed.",
        "type": "string"
    }
},
"required": ["time", "host", "source", "sourcetype", "index"]
}

```

GeneralSplunkEvent.json

```

{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "systemUniqueId": {
      "description": "The UUID of the system.",
      "type": "string"
    },
    "systemVersion": {
      "description": "The release version of the system.",
      "type": "string"
    }
  },
  "required": ["systemUniqueId", "systemVersion"]
}

```

Action.json

```

{
  "$schema": "http://json-schema.org/draft-04/schema#",

```

```

"definitions": {
  "ActionEvent": {
    "type": "object",
    "properties": {
      "reference": {
        "description": "The object reference of the action.",
        "type": "string"
      },
      "title": {
        "description": "Action title.",
        "type": "string"
      },
      "details": {
        "description": "Plain text description of the action.",
        "type": "string"
      },
      "startTime": {
        "description": "The time the action occurred. For a long running
process, this represents the starting time.",
        "type": "string"
      },
      "endTime": {
        "description": "The time the action completed.",
        "type": "string"
      },
      "user": {
        "description": "The user who initiated the action.",
        "type": "string"
      },
      "userAgent": {
        "description": "Name of the client software used to initiate the
action.",
        "type": "string"
      },
      "originIp": {
        "description": "Network address used to initiate the action",
        "type": "string"
      },
      "parentAction": {
        "description": "The parent action of this action.",
        "type": "string"
      },
      "state": {
        "description": "State of the action",
        "type": "string"
      },
      "workSource": {
        "description": "Origin of the work that caused the action.",
        "type": "string"
      },
      "workSourceName": {

```

```

        "description": "Name of the user or policy that initiated the
action.",
        "type": "string"
    },
    "failureDescription": {
        "description": "Details of the action failure.",
        "type": "string"
    },
    "failureAction": {
        "description": "Action to be taken to resolve the failure",
        "type": "string"
    },
    "failureMessageCode": {
        "description": "Message ID associated with the event.",
        "type": "string"
    }
},
    "required": ["reference", "title", "details", "state"]
}
},
"type": "object",
"title": "Action",
"allOf": [{ "$ref": "GeneralSplunkHeader.json#" }],
"properties": {
    "event": {
        "type": "object",
        "allOf": [{ "$ref": "#/definitions/ActionEvent" }, { "$ref":
GeneralSplunkEvent.json#" }]
    }
}
}
}

```

Alert.json

```

{
    "$schema": "http://json-schema.org/draft-04/schema#",
    "definitions": {
        "AlertEvent": {
            "type": "object",
            "properties": {
                "reference": {
                    "description": "The object reference of the alert.",
                    "type": "string"
                },
            },
            "title": {
                "description": "Title of the event which triggered the alert.",
                "type": "string"
            },
        },
        "code": {

```

```

        "description": "Dotted descriptor of the type of event which
triggered the alert.",
        "type": "string"
    },
    "eventSeverity": {
        "description": "The severity of the event.",
        "type": "string"
    },
    "details": {
        "description": "Plain text description of the event which
triggered the alert.",
        "type": "string"
    },
    "response": {
        "description": "Automated response, if any, taken by the system.",
        "type": "string"
    },
    "timestamp": {
        "description": "The time the alert occurred.",
        "type": "string"
    },
    "target": {
        "description": "Reference to the target object.",
        "type": "string"
    },
    "targetName": {
        "description": "Name of the target object.",
        "type": "string"
    }
    },
    "required": ["reference", "title", "code", "details", "timestamp"]
}
},
"type": "object",
"title": "Alert",
"allOf": [{ "$ref": "GeneralSplunkHeader.json#" }],
"properties": {
    "event": {
        "type": "object",
        "allOf": [{ "$ref": "#/definitions/AlertEvent" }, { "$ref":
GeneralSplunkEvent.json#" }]
    }
}
}
}

```

JobEvent.json

```

{
    "$schema": "http://json-schema.org/draft-04/schema#",

```

```

"definitions": {
  "JobEvent": {
    "type": "object",
    "properties": {
      "job": {
        "description": "The object reference of job associated with this
event.",
        "type": "string"
      },
      "parentAction": {
        "description": "The object reference of the parent action of the
associated job.",
        "type": "string"
      },
      "jobState": {
        "description": "The new state of the job.",
        "type": "string"
      },
      "timestamp": {
        "description": "The time the event occurred.",
        "type": "string"
      },
      "percentComplete": {
        "description": "Completion percentage of the job associated with
this event",
        "type": "number"
      },
      "diagnoses": {
        "description": "If job failed, a set of diagnoses of things that
may have caused the failure.",
        "type": "array",
        "items": { "$ref": "#/definitions/DiagnosisResult" }
      },
      "eventType": {
        "description": "The type of this event (info, warning, or
error).",
        "type": "string"
      },
      "messageCode": {
        "description": "A message code describing this event.",
        "type": "string"
      },
      "messageDetails": {
        "description": "A message describing the details of this event.",
        "type": "string"
      },
      "messageAction": {
        "description": "Action to be taken by the user to repair or
remedy the situation.",
        "type": "string"
      },
      "messageCommandOutput": {

```

```

        "description": "Any command output generated by this event",
        "type": "string"
    },
    "required": ["job", "parentAction", "timestamp", "percentComplete",
"messageCode", "messageDetails", "eventType"]
},
"DiagnosisResult": {
    "type": "object",
    "properties": {
        "diagnosisCode": {
            "description": "Message code associated with this diagnosis
check.",
            "type": "string"
        },
        "diagnosisMessage": {
            "description": "Description of this diagnosis check.",
            "type": "string"
        },
        "failed": {
            "description": "True if this diagnosis check did not pass.",
            "type": "boolean"
        },
        "targetReference": {
            "description": "Reference of the target object of this diagnosis
check, if applicable.",
            "type": "string"
        }
    },
    "required": ["diagnosisCode", "diagnosisMessage", "failed"]
},
"DiagnosisResult": {
    "type": "object",
    "title": "JobEvent",
    "allOf": [{ "$ref": "GeneralSplunkHeader.json#" }],
    "properties": {
        "event": {
            "type": "object",
            "allOf": [{ "$ref": "#/definitions/JobEvent" }, { "$ref":
"GeneralSplunkEvent.json#" } ]
        }
    }
}
}
}

```

AbstractFault.json

```

{
    "$schema": "http://json-schema.org/draft-04/schema#",
    "type": "object",
    "title": "AbstractFault",

```

```

"properties": {
  "reference": {
    "description": "The object reference of the fault.",
    "type": "string"
  },
  "title": {
    "description": "Title of the event which triggered the fault.",
    "type": "string"
  },
  "code": {
    "description": "Dotted descriptor of the type of event which triggered
the fault.",
    "type": "string"
  },
  "details": {
    "description": "Plain text description of the event which triggered the
fault.",
    "type": "string"
  },
  "response": {
    "description": "Automated response, if any, taken by the system.",
    "type": "string"
  },
  "dateDiagnosed": {
    "description": "The date when the fault was diagnosed.",
    "type": "string"
  },
  "target": {
    "description": "Reference to the target object.",
    "type": "string"
  },
  "targetName": {
    "description": "Name of the target object.",
    "type": "string"
  },
  "state": {
    "description": "The state of the fault."
  },
  "eventSeverity": {
    "description": "The severity of the event.",
    "type": "string"
  }
},
"required": ["reference", "title", "details", "state", "target", "dateDiagnosed"]
}

```

Fault.json

```
{
```

```

"$schema": "http://json-schema.org/draft-04/schema#",
"definitions": {
  "FaultEvent": {
    "type": "object",
    "properties": {
      "dateResolved": {
        "description": "The date when the fault was resolved.",
        "type": "string"
      },
      "resolutionComments": {
        "description": "Comments regarding the resolution of the fault.",
        "type": "string"
      }
    }
  }
},

"type": "object",
"title": "Fault",
"allOf": [{ "$ref": "GeneralSplunkHeader.json#" }],
"properties": {
  "event": {
    "type": "object",
    "allOf": [{ "$ref": "AbstractFault.json#" }, { "$ref": "#/definitions/
FaultEvent" }, { "$ref": "GeneralSplunkEvent.json#" } ]
  }
}
}

```

FaultEffect.json

```

{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "definitions": {
    "AlertEvent": {
      "type": "object",
      "properties": {
        "reference": {
          "description": "The object reference of the alert.",
          "type": "string"
        },
        "title": {
          "description": "Title of the event which triggered the alert.",
          "type": "string"
        },
        "code": {
          "description": "Dotted descriptor of the type of event which
triggered the alert.",
          "type": "string"
        }
      }
    },

```

```

        "eventSeverity": {
            "description": "The severity of the event.",
            "type": "string"
        },
        "details": {
            "description": "Plain text description of the event which
triggered the alert.",
            "type": "string"
        },
        "response": {
            "description": "Automated response, if any, taken by the system.",
            "type": "string"
        },
        "timestamp": {
            "description": "The time the alert occurred.",
            "type": "string"
        },
        "target": {
            "description": "Reference to the target object.",
            "type": "string"
        },
        "targetName": {
            "description": "Name of the target object.",
            "type": "string"
        }
    },
    "required": ["reference", "title", "code", "details", "timestamp"]
},
"type": "object",
"title": "Alert",
"allOf": [{ "$ref": "GeneralSplunkHeader.json#" }],
"properties": {
    "event": {
        "type": "object",
        "allOf": [{ "$ref": "#/definitions/AlertEvent" }, { "$ref":
"GeneralSplunkEvent.json#" } ]
    }
}
}

```

Types of Events

Delphix uses the **source** field in Splunk to designate the type of each event. Here is a full list of the possible values of the **source** field for events, along with an explanation of when each event is generated and the name of the corresponding JSON schema that describes the event structure.

source	Explanation	Schema
delphix.events.action.started	Action has started running.	Action.json

source	Explanation	Schema
delphix.events.action.waiting	Action has moved to the WAITING state.	Action.json
delphix.events.action.completed	Action has completed successfully.	Action.json
delphix.events.action.failed	Action has failed.	Action.json
delphix.events.action.canceled	Action has been canceled.	Action.json
delphix.events.job.event	Job Event has been generated in response to a Job progress update.	JobEvent.json
delphix.events.fault.posted	Fault has been posted.	Fault.json
delphix.events.fault.resolved	Fault has been resolved.	Fault.json
delphix.events.fault.ignored	User has chosen to ignore a fault.	Fault.json
delphix.events.fault.unignored	User has chosen to “unignore” a previously ignored fault.	Fault.json
delphix.events.fault.effect.posted	Fault Effect has been posted as a downstream effect of some Fault.	FaultEffect.json
delphix.events.fault.effect.resolved	Fault Effect has been resolved as a result of its cause being resolved.	FaultEffect.json
delphix.events.fault.effect.ignored	Fault Effect has been ignored when a user chose to ignore its cause.	FaultEffect.json
delphix.events.fault.effect.unignored	Fault Effect has been “unignored” when a user chose to unignore its previously ignored cause.	FaultEffect.json
delphix.events.alert	Alert has been posted.	Alert.json

Search Examples - Metric Formats

JSON Schemas for Metrics

GeneralMetricEvent.json

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "type": "object",
  "properties": {
    "systemUniqueId": {
      "description": "The UUID of the system.",
      "type": "string"
    },
    "systemVersion": {
      "description": "The release version of the system.",
      "type": "string"
    },
    "event": {
      "description": "A tag that indicates this Event is a metric.",
      "type": "string",
      "enum": ["metric"]
    }
  },
  "required": ["systemUniqueId", "systemVersion", "event"]
}
```

A schema that is common for all metrics, the “data” nested JSON object contents varies depending on the metric:

CommonMetric.json

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "definitions": {
    "CommonMetric": {
      "type": "object",
      "properties": {
        "name": {
          "description": "The name of the metric.",
          "type": "string"
        },
        "time": {
          "description": "The timestamp of the metric when it was collected.  
The default time format is epoch time format, in the format <sec>.<ms>.",
          "type": "number"
        },
        "value": {
          "description": "The numeric value of the metric.",
          "type": "number"
        }
      }
    }
  }
}
```


JSON key	Value Type	Description	Example	Comments
sourcetype	"_json"	Data format	n/a	Used for Splunk Indexed field extractions
systemUniqueId	String	The UUID of the system	"423f22db-4ee9-6ebe-ff0f-884ffdc351f7"	
systemVersion	String	The release version of the system	5.2.5.0	

And the general key-values specific to an Insight metric are:

JSON key	Value Type	Description	Example	Comments
name	String	Pseudo hierarchical dotted format of the metric name.	system.cpu.util.pct system.disk.ops.co unt system.net .total.bytes	A metric name: - is alphanumeric (underscore & dot allowed) - has <i>prefix</i> that points to the source of the data (like <i>system.cpu</i>) - has <i>suffix</i> that describes the unit (when it can) - an aggregate/summary metric would have <i>total</i> as part of suffix for example.
time	UNIX time notation (epoch)	Timestamp	1525399950	UNIX epoch time format.
value	Numeric value of the metric	The actual measurement	85.17 (system.cpu.util.pct)	Numeric only
type	String (Enum): - counter, - gauge, - summary	Type of the measurement	counter gauge summary	This field could be used to identify what kind of value a metric is presenting (like a "gauge" for cpu metric implies the values will fall into a range of 0-100% - similarly a "summary" would imply the value is an "accumulated value" like network utilization)

JSON key	Value Type	Description	Example	Comments
data	JSON Object	Dimensions of the metric (if any).	<pre>{ "read_latency": 0.98, "ops_write": 20, "ops_read": 30, "write_latency": 0.55, }</pre>	<p>Has additional info about a metric (called “Dimensions” in Splunk)</p> <p>Has different key-vals depending on a given metric</p> <p>Like for CPU: it can have key-vals to like “kernel”, “user”</p> <p>For disk: can have latency/throughput, operation name</p>

Summary of all the current Metrics

Metric name	Metric value	Metric dimensions (data.xxx below)
system.cpu.util.pct	cpu utilization percentage aggregated across all cpus	user, kernel, idle
system.disk.ops.count	disk read/write operations aggregated across all disk instances	count, op (read/write), latency, avgLatency, throughput
system.net.total.bytes	network utilization - total bytes (in + out) of a given network interface	networkInterface, inBytes, outBytes, inPackets, outPackets
system.nfs.ops.count	nfs read/write operations aggregated across all instances	count, op (read/write), latency, avgLatency, throughput
system.iscsi.ops.count	iscsi read/write operations aggregated across all instances	count, op (read/write), latency, avgLatency, throughput
system.tcp.total.bytes	tcp connection statistics aggregated by service (nfs/iscsi/dsp etc)	congestionWindowSize, inBytes, inUnorderedBytes, localAddress, outBytes, receiveWindowSize, remoteAddress, retransmittedBytes, roundTripTime, sendWindowSize, service unacknowledgedBytes, unsentBytes
system.dataset.total.bytes	dataset performance - total number of bytes read + written per dataset (dsource/vdb etc)	dataset, nread, nwritten, type (virtual, dsource, staging)

Metric name	Metric value	Metric dimensions (data.xxx below)
<i>system.capacity.source.size</i>	overall system capacity - actual used space in bytes	actualSpace, totalSpace, activeSpace, actualSpace, descendantSpace, logSpace, manualSpace, policySpace, syncSpace, timeflowUnvirtualizedSpace, unownedSnapshotSpace, unvirtualizedSpace
<i>system.capacity.consumer.size</i>	consumer capacity - actual used space in bytes per consumer (vdb/pdb etc)	syncSpace, descendantSpace, activeSpace, timeflowUnvirtualizedSpace, objectType, dSource, actualSpace, groupName, manualSpace, unownedSnapshotSpace, unvirtualizedSpace, logSpace, policySpace, consumerName

Related Topics

- [Splunk Integration](#)
- [Configuring Splunk](#)

Fluentd Plugin Service for API Modules

Overview

Beginning with 5.3.0.0 release, Delphix engines have had the ability to forward events and metrics to a Splunk host using Fluentd. In 6.0.13.0 Delphix has extended this capability for use with other monitoring packages (ELK Stack, Datadog, etc.). Forwarding to Splunk can still be configured out of the box by providing about the Splunk host and other parameters. Fluentd plugins provide a mechanism to customize the Delphix engine's forwarding capabilities for output to other data consumers.


The Fluentd plugin is a tar file with two components. A fluentd configuration file template is required. Optionally, any gem files that are needed to execute the directives in the configuration file can be included. Once the plugin is uploaded, the configuration occurs by collecting user input for the variables required by the template. The user interfaces adapts based on the variables required by the configuration template.

The current implementation is limited to one data consumer at a time. Only one plugin, in addition to the default SplunkHec plugin, can be uploaded. This is an intended simplification for the initial Fluentd release version.

i The current implementation is limited to one data consumer at a time. Only one plugin, in addition to the default SplunkHec plugin, can be uploaded. This is an intended simplification for the initial Fluentd release version.

Technical Details

Architectural diagram



GUI, API or CLI changes (if any)

APIs with a corresponding CLI for uploading plugins and Fluentd configuration have been added. The service/fluentd/plugins API displays the plugins that are available.

```

service fluentd plugins> ls

Objects      REFERENCE          PLUGIN          GEMS  ATTRIBUTEDEFINITIONS
SCHEMADEFINITION

FLUENTD_PLUGIN-1      splunkHec  ...           ...
...

Operations

requestUploadToken

```

The splunkHec plugin appears by default and cannot be deleted. Additionally, the requestUploadToken operation provides a token needed to upload a plugin via the data/upload API. If a second plugin is uploaded, it will appear in the list. Before a third plugin could be uploaded, the second must be selected and deleted.

The following curl commands illustrate the use of the API to upload a plugin

```

curl -s -X POST -k --data @- http://<engine-name>/resources/json/delphix/session -c
~/cookies.txt -H "Content-Type: application/json" <<EOF

{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 4,
    "micro": 3
  }
}

```

```
EOF
```

```
curl -s -X POST -k --data @- http://<engine-name>/resources/json/delphix/login -b ~/cookies.txt -c ~/cookies2.txt -H "Content-Type: application/json" <<EOF
```

```
{
  "type": "LoginRequest",
  "username": "sysadmin",
  "password": "sysadmin",
  "target": "SYSTEM"
}
```

```
EOF
```

```
curl -s -X POST -k --data @- http://<engine-name>/resources/json/delphix/service/fluentd/plugins/requestUploadToken -b ~/cookies2.txt -H "Content-Type: application/json" <<EOF
```

```
Returns the token e.g: {"type":"OKResult","status":"OK","result":
{"type":"FileUploadResult","url":"/resources/json/delphix/data/
upload","token":"59346df5-3ecd-4ced-afbf-97acefa156dc"},"job":null,"action":null}
```

```
curl -b ~/cookies2.txt -X POST -F "file=@/Users/blewis/ws/far-dev/
splunk_host_port.far" -F "token=<token>" http://<engine-name>/resources/json/delphix/
data/upload
```

The service/fluentd/plugins API displays the tokens that are available.

An example of creating a configuration with the default splunkHec plugin

```
service fluentd configuration> create

service fluentd configuration create *> set plugin=splunkHec; set enabled=true; edit
attributes

service fluentd configuration create attributes *> add; set name=hec_host; set
value=http://vmname-splunkhost.delphix.com; back

service fluentd configuration create attributes *> add; set name=hec_port; set
value=8088; back;

service fluentd configuration create attributes *> add; set
type=FluentdSecretAttribute;

service fluentd configuration create attributes 2 *> set name=hec_token; set
secretValue=bb75c032-bdea-4c19-b152-d158cf13e019; back
```

```
service fluentd configuration create attributes *> add; set name=eventsIndex; set value=delphix_events; back;
```

```
service fluentd configuration create attributes *> add; set name=protocol; set value=https; back
```

```
service fluentd configuration create attributes *> add; set name=metricsPushFrequency; set value=60; back
```

```
service fluentd configuration create attributes *> add; set name=metricsIndex; set value=delphix_metrics; back
```

```
service fluentd configuration create attributes *> add; set name=eventsPushFrequency; set value=60; back;
```

```
service fluentd configuration create attributes *> ls
```

Properties

0:

type: FluentdRegularAttribute (*)

name: hec_host (*)

value: http://vmname-splunkhost.delphix.com (*)

1:

type: FluentdRegularAttribute (*)

name: hec_port (*)

value: 8088 (*)

2:

type: FluentdSecretAttribute (*)

name: hec_token (*)

secretValue: ***** (*)

3:

type: FluentdRegularAttribute (*)

name: eventsIndex (*)

value: delphix_events (*)

```
4:
  type: FluentdRegularAttribute (*)
  name: protocol (*)
  value: https (*)
```

```
5:
  type: FluentdRegularAttribute (*)
  name: metricsPushFrequency (*)
  value: 60 (*)
```

```
6:
  type: FluentdRegularAttribute (*)
  name: metricsIndex (*)
  value: delphix_metrics (*)
```

```
7:
  type: FluentdRegularAttribute (*)
  name: eventsPushFrequency (*)
  value: 60 (*)
```

Use the "add" command to add an element to this array.

```
service fluentd configuration create attributes *> back
```

```
service fluentd configuration create *> ls
```


Properties

```
type: FluentdConfig
name: (unset)
attributes: [ ... ] (*)
enabled: true (*)
plugin: splunkHec (*)
```


```
service fluentd configuration create *> commit
```

Creating Default splunkHec Configuration From GUI


1. Login as a sysadmin user and select **Fluentd Configuration** under the Preferences menu.




2. The splunkHec plugin appears in the dropdown by default.



- 3. Additional plugins can be added by using the + button on the right.



- 4. The new plugin uploaded can be viewed in the dropdown menu and selected to add a configuration.



5. The delete button on the right can be used to delete a selected plugin.
6. The name of fields for configuration would be the same as those provided in the uploaded plugin file.

The screenshot shows a 'Fluentd Configuration' window with a close button (X) in the top right corner. Below the title bar, there is a dropdown menu labeled 'Select a plugin Configuration' with 'splunkHec' selected, and buttons for adding (+) and deleting (trash) configurations. The configuration fields are as follows:

Field Name	Value
protocol	https
hec_port	8000
metricsPushFrequency	60
metricsIndex	dlpx_metrics
eventsIndex	dlpx_events
eventsPushFrequency	5
hec_host	qa-splunk.ops.delphix.com
hec_token	d64ee2e5-2fd7-46dd-8029-13544ad7ee92

At the bottom right, there are 'Close' and 'Save' buttons.

7. The configuration can be saved by clicking the save button on the bottom.
8. When loading an existing configuration on UI, the secret fields are masked.

Implementation

Secret attributes have been introduced in order to protect private data such as passwords and tokens. In the CLI and API, they are identified by the **FluentdSecretAttribute** type. In plugins, attribute names in the config file template beginning with **_secret** indicate secret data.

Known issues

Fluentd runs in a systemd container to limit negative side-effects from a bad plugin. More security can be added to this container by running as a non-root user, in addition to limiting the size of the log and buffer files it can write.

Performance Analytics Management

Delphix offers various performance analytics tools to help users monitor throughput, latency, and other key metrics. Learn more about how to leverage these tools and how to architect your Delphix deployment for optimal performance.

This section covers the following topics:

- [Performance Analytics](#)
- [Storage Performance Configuration Options](#)
- [Architecture for Performance - Hypervisors and Host](#)
- [Target Host OS and Database Configuration Options](#)

Performance Analytics

This section covers the following topics:

- [Performance Analytics Tool Overview](#)
- [Working with Performance Analytics Graphs in the Graphical User Interface](#)
- [Performance Analytics Statistics Reference](#)
- [Performance Analytics Tool API Reference](#)
- [Performance Analytics Case Study: Using a Single Statistic](#)
- [Performance Analytics Case Study: Using Multiple Statistics](#)

Performance Analytics Tool Overview

This topic describes the Performance Analytics tool and illustrates some basic uses of it.

Introduction

The performance analytics tool allows introspection into how the Delphix Engine is performing. The introspection techniques it provides are tuned to allow an iterative investigation process, helping to narrow down the cause associated with the performance being measured. Performance analytics information can be accessed through the Delphix Management application, as described in [Working with Performance Analytics Graphs in the Graphical User Interface](#), as well as the CLI and the web services API, as described in other topics in this section. The default statistics that are being collected on the Delphix Engine include CPU utilization, network utilization, and disk, NFS, and iSCSI IO operations (see [Performance Analytics Statistics Reference](#) for details).

The performance tool operates with two central concepts: **statistics** and **statistic slices**.


Statistics

Each statistic describes some data that can be collected from the Delphix Engine. The first piece of information a statistic provides is its **type**, which you will use as a handle when creating a statistic slice. It also gives the **minimum collection interval**, which puts an upper bound on the frequency of data collection. The actual data a statistic can collect is described through a set of **axes**, each of which describes one "dimension" of that statistic. For example, the statistic associated with Network File System (NFS) operations has a **latency axis**, as well as an **operation type axis** (among many others), which allows users to see NFS latencies split by whether they were reads or writes.


Each axis has some important information embedded in it.

- The **name** of the axis provides a short description of what the axis collects and is used when creating a statistic slice
- A **value** type, which tells you what kind of data will be collected for this axis. The different value types are **integer**, **boolean**, **string**, and **histogram**. The first three are straightforward, but statistic axes with a histogram type can collect a distribution of all the values encountered during each collection interval. This means that instead of seeing an average NFS operation latency every collection interval, you can see a full distribution of operation latencies during that interval. This allows you to see outliers as well as the average, and observe the effects of caching on the performance of your system more easily.
- A **constraint** type, which is only relevant while creating a statistic slice, and will be described in more detail below

One last bit of information that an axis provides makes the most sense after seeing how data points are queried. In the most basic situation, you would only collect one axis of a statistic, such as the latency axis from the NFS operations statistic. When you ask for data, you would get back a data point for every collection interval in the time range you requested. These data points would be grouped into a single stream.



However, if you had collected the operation type axis as well as the latency axis, you would get two streams of datapoints: one for reading operations, and one for write operations.



Because the operation axis applies to many data points, the data points returned are split into two streams, and the operation axis is stored with the top-level stream instead of with each data point in the streams. However, the latency axis will be different for each data point in a stream, so it is not an attribute of the stream, but instead an attribute of the datapoint.

Statistic Slices

Statistics describe what data can be collected and are auto-populated by the system, but statistic slices are responsible for actually collecting the data, and you must create them manually when you want to collect some performance data. Each slice is an instantiation of exactly one statistic, and can only gather data that is described by that statistic. "Slices" are so named because each one provides a subset of the information available from the parent statistic it is associated with. A statistic can be thought of as describing the axes of a multidimensional

space, whereas you typically will only want to collect a simpler slice of that space due to the large number of axes available.

When you specify a slice, there are several fields that you must supply:

- The statistic type this slice is associated with. This must be the same type as the statistic of which this is an instantiation.
- The collection interval, which must be greater than the minimum collection interval the parent statistic gives
- The axes of the parent statistic this slice will collect

Finally, a slice can place constraints on axes of its parent statistic, allowing you to limit the data you get back. For instance, if you're trying to narrow down the cause of some high NFS latency outliers, it may be useful to filter out any NFS latencies which are shorter than one second. To do this, you would place a constraint on the latency axis of an NFS operation slice that states that the values must be higher than one second. You can constrain any axis in the same fashion, and each axis' description in the parent statistic gives a constraint type that can be applied to it. This allows you to place different types of constraints on the latency axis (which is a number measured in nanoseconds) than the operation type axis (which is an enum that can take the values "read" or "write").

Persisting Analytics Data

Data collected by slices is persisted temporarily on the Delphix Engine. Performance data is guaranteed to be available at the finest resolution for six hours, then is compressed to per-minute data and held for seven days, and finally compressed to per-hour data and held for 30 days. If data of a certain resolution will be needed longer than these limits, you should instruct the slice to remember the data permanently until you are done using it. The commands to manage this are listed in the [Performance Analytics Tool API Reference](#).

Related Topics

- [Working with Performance Analytics Graphs in the Graphical User Interface](#)
- [Performance Analytics Tool API Reference](#)
- [Performance Analytics Statistics Reference](#)

Working with Performance Analytics Graphs in the Graphical User Interface

This topic describes the performance analytics graphs that are available in the Delphix Management application, and the controls for changing the views of those graphs.

Accessing the Performance Analytics Graphs

1. Log into the Delphix Management application.
2. In the **Resources** menu, select **Performance Analytics**.
3. Use the controls described below to view statistics and their related graphs.

General Graph Display and Controls

Control Name	Control Description	Usage
Chart Tabs	Specifies which graphs/charts are displayed.	Chart tabs have a multiple select drop-down which has the same three options checked by default (CPU, Network, Disk IO). Users can show/hide graphs by checking/unchecking drop-down options.
Timescale	Controls the time range of data displayed in the graph. Available values are 1 minute, 1 hour, and 1 day. By default, 1 minute is selected.	Timescale options are presented as a drop-down, select from Minute , Hour , or Day to change the Zoom Level.
Pause/Resume	Pause/Resume icons	All charts display live data so users can pause it by clicking the Pause button in the upper right corner of each graph. Once the chart is paused, the Pause button is replaced by the Resume button so a user can click it to see live data again.
Op. Type	Operation Type (used in Disk IO, NFS and iSCSI charts)	Options are presented as a drop-down.
Shown Data Timeline	Displays timestamps of data points in the graph.	
Available Data Timeline	Displays navigable time ranges for historical data.	

Control Name	Control Description	Usage
Timeline Selector	<p>Specifies the start and end time for the currently displayed data. The range displayed is controlled by the Zoom Level.</p>	<p>Drag the Timeline Selector to view statistics for a specific time in the past, or click the scroll bar arrows to view the desired time period. You can also use the slider controls within the Timeline Selector to change the length of time for which data is displayed.</p> <p>When the Timeline Selector is aligned to the right of the timeline, it represents live data that is updated every second. If the Timeline Selector is moved from the right alignment with the timeline, the data displayed is historical and no live updates are displayed. To resume live data updates, move the Timeline Selector back to the right-aligned position representing the current time. The data will be refreshed to the latest data, and live updates will resume every second.</p>
Graph Legend	<p>If more than one set of information is presented on the graph, the Graph Legend displays a description and color for each set and allows a user to toggle that set-off and on.</p> <p>For example, in the network graph there can be multiple network interfaces, and for each network interface the graph displays four statistics (bytes sent, bytes received, packets sent, packets received). When a user toggles off a specific network interface, all four statistics corresponding to that interface are hidden from the screen.</p> <p>The color for lines representing bytes sent and packets sent is the same. Similarly, the color for lines representing bytes received and packets received is the same. This makes it easier to correlate the number of bytes and the number of packets sent/received for a given network interface.</p>	<p>To hide a set of information, click on the set name within the Graph Legend. Data representing that set is removed from the graph, and the set's name is greyed out. To show a set that has been hidden, click on the set name.</p>

Latency, Timeline Page, and Tooltip Graph Display and Controls

Control Name	Control Description	Usage
Timeline Page Left/ Right Button	Scrolls Available Data Timeline by a specified time range depending on the current Zoom Level .	When the Zoom Level is set to Minute , click Timeline Page Left . The Available Data Timeline is changed to show the time period for the previous hour prior.
Graph Value Tooltip	Shows a value, along with the time stamp, for a specific data point.	Mouse over a data point on the graph to view the tooltip.
Latency Range Selector (shown on latency heatmaps only)	Controls the lower and upper limits for displayed latency buckets.	Drag the lower and upper controls to drill down into a specific range of latency buckets. Latency buckets that fall outside of the selected range are summarized, the lower row representing latency buckets that are below the lower limit, and the upper row representing latency buckets that are below the upper limit of the latency range selector. Use Latency Range Selector to view a more detailed distribution of latencies for a specific range.
Latency Outlier Selector (shown on latency heatmaps only)	Hides infrequent latencies (outliers) based on a percentage threshold. Its range is 0%-10%, with a default of 0%. The percentage establishes a threshold below which buckets are considered "outliers" and are hidden from the graph. Each bucket is assigned a percentage based on the ratio of its count vs the maximum count of any bucket in the graph.	Drag the control to the desired percentage threshold.

Related Topics

- [Performance Analytics Tool Overview](#)
- [Performance Analytics Statistics Reference](#)

Performance Analytics Statistics Reference

This topic describes the various performance statistics that are available for the Delphix Engine and how they can be used to analyze and improve performance.

The Delphix Engine is shipped with a default set of statistics that are collected on Delphix Engine virtual appliance, as listed below. The statistics are stored for up to 30 days for historical analysis.

Statistic	Description
CPU Utilization	Total CPU utilization for all CPUs. This statistic includes both kernel and user time.
Network Throughput	Measures throughput in bytes and packets, broken down by sent vs. received data and by the network interface. Each network interface shows four graphed lines: bytes sent, bytes received, packets sent, and packets received. To help easily correlate bytes and packets, the same color is used for both bytes and packet values.
Disk IO	Measures a number of IO operations, and the latencies and throughput of the underlying storage layer. The statistic is represented by the graphs - a column chart for IO operations, a heat map for latency distribution, and a line chart for throughput. IO operations are grouped by reads and write. A shaded rectangle on a latency heat map represents an IO operation (read or write) which falls within a particular time range (bucket). The shading of rectangles depends on the number of IO operations that fall within a particular bucket - the higher the count the darker the shading.
NFS	Measures a number of IO operations and the latencies and throughput of the NFS server layer in the Delphix Engine. Its graphical representation is similar to the Disk IO graph. It is useful to diagnose the performance of dSources and VDBs that use NFS mounts (Oracle, PostgreSQL).
iSCSI	Measures the number of IO operations, and the latencies and throughput, of the iSCSI server layer in the Delphix Engine. Its graphical representation is similar to the Disk IO graph. It is useful to diagnose the performance of Microsoft SQL Server dSources and VDBs.

Related Topics

- [Performance Analytics Tool Overview](#)
- [Performance Analytics Tool API Reference](#)

Performance Analytics Tool API Reference

This topic describes basic commands and command syntax for using the Performance Analytics tool.

Statistic Types

More detailed information can be found about each statistic type through the command-line interface (CLI) and webservice API, but the following table provides more information about how similar I/O stack statistic types relate to each other.

Statistic Type	Description	Axis Name	Axis Description	Axis Value Type
NFS_OPS	Provides information about Network File System operations. This is the entry point to the Delphix Engine for all Oracle database file accesses.	op	I/O operation type	STRING
		path	Path of the affected file	STRING
		size	I/O sizes in bytes	HISTOGRAM
		avgLatency	Average I/O latency in nanoseconds	INTEGER
		cached	Whether reads were cached	BOOLEAN
		latency	I/O latencies in nanoseconds	HISTOGRAM
		count	Number of I/O operations	INTEGER
		client	Address of the client	STRING
		throughput	I/O throughput in bytes	INTEGER
		sync	Whether writes were synchronous	BOOLEAN
iSCSI_OPS	Provides information about iSCSI operations. This is the entry point to the Delphix Engine for all SQL Server file accesses.	Same axes as NFS_OPS, except for path, cached, and sync.		

Statistic Type	Description	Axis Name	Axis Description	Axis Value Type
VFS_OPS	This layer sits immediately below NFS_OPS and iSCSI_OPS. It should give almost exactly the same latencies, assuming no unexpected behavior is occurring.	Same axes as NFS_OPS, except for client.		
DxFS_OPS	This layer sits immediately below VFS_OPS, and the two of them should give almost exactly the same latencies.	Same axes as VFS_OPS.		
DxFS_IO_QUEUE_OPS	This layer sits below DxFS_OPS, but the latencies will differ from that layer because this layer batches together operations to increase throughput.	op	I/O operation type	STRING
		count	Number of I/O operations	INTEGER
		size	I/O sizes in bytes	HISTOGRAM
		avgLatency	Average I/O latency in nanoseconds	INTEGER
		latency	I/O latencies in nanoseconds	HISTOGRAM
		throughput	I/O throughput in bytes	INTEGER
		priority	Priority of the I/O	STRING
DISK_OPS	This layer sits below DxFS_IO_QUEUE_OPS at the bottom of the I/O stack and measures interactions between the Delphix Engine and disks.	op	I/O operation type	STRING
		count	Number of I/O operations	INTEGER
		size	I/O sizes in bytes	HISTOGRAM
		avgLatency	Average I/O latency in nanoseconds	INTEGER

Statistic Type	Description	Axis Name	Axis Description	Axis Value Type
		latency	I/O latencies in nanoseconds	HISTOGRAM
		throughput	I/O throughput in bytes	INTEGER
		error	Whether the I/O resulted in an error	BOOLEAN
		device	Device the I/O was issued to	STRING
CPU_UTIL	This is unrelated to the layers of the I/O stack. It measures CPU utilization on the Delphix Engine.	idle	Idle time in milliseconds (showAxes command may incorrectly state nanoseconds)	INTEGER
		user	User time in milliseconds (showAxes command may incorrectly state nanoseconds)	INTEGER
		kernel	Kernel time in milliseconds (showAxes command may incorrectly state nanoseconds)	INTEGER
		dtrace	DTrace time in milliseconds (showAxes command may incorrectly state nanoseconds) Subset of time in kernel	INTEGER
		cpu	Which CPU was utilized	INTEGER
NETWORK_INTERFACE_UTIL	Network interface utilization on the Delphix Engine.	inBytes	Number of bytes received	INTEGER
		inPackets	Number of packets received	INTEGER

Statistic Type	Description	Axis Name	Axis Description	Axis Value Type
		outBytes	Number of bytes transmitted	INTEGER
		outPackets	Number of packets transmitted	INTEGER
		networkInterface	Which network interface was utilized	STRING
TCP_STATS	Statistics for all established TCP connections on the Delphix Engine.	localAddress	Local address for the TCP connection	STRING
		localPort	Local port for the TCP connection	INTEGER
		remoteAddress	Remote address for the TCP connection	STRING
		remotePort	Remote port for the TCP connection	INTEGER
		inBytes	Data bytes received	INTEGER
		outBytes	Data bytes transmitted	INTEGER
		receiveWindowSize	The size of the local receive window	INTEGER
		sendWindowSize	The size of the peer's receive window	INTEGER
		congestionWindowSize	The size of the local congestion window	INTEGER
		retransmittedBytes	Bytes retransmitted	INTEGER
		inUnorderedBytes	Number of bytes received out of order. This is a subset of the "inBytes" value	INTEGER

Statistic Type	Description	Axis Name	Axis Description	Axis Value Type
		unacknowledgedBytes	Number of bytes sent but unacknowledged	INTEGER
		roundTripTime	Smoothed average round-trip time in microseconds	INTEGER

Statistic Axis Value Types

Values are returned when a slice's data is queried. Each axis has a value type, which specifies how the data will be returned.

Value Type	Description
INTEGER	The value is returned as an integer. For information about what units the integer is measured in, read the documentation for the related datapoint or datapoint stream type.
BOOLEAN	The value is returned as a boolean.
STRING	The value is returned as a string. This is used for enum values as well, although the set of strings that can be returned is limited.
HISTOGRAM	<p>The value is returned as a log-scale histogram. The histogram has size buckets whose minimum and maximum value get doubled. Histograms are returned as JSON maps, where the keys are the minimum value in a bucket and the values are the height of each bucket.</p> <p>Here is an example histogram. Notice that buckets with a height of zero are not included in the JSON object and that keys and values are represented as strings.</p> <pre> { "32768": "10", "65536": "102", "262144": "15", "524288": "2" } </pre>

Axis constraints are used to limit the data which a slice can collect. Each axis specifies a constraint type which can be used to limit that axis' values.

Constraint Type	Description
Constraint Type	Description
<code>BooleanConstraint</code>	A superclass in which constraints on boolean values must extend. Currently, the only subclass is <code>BooleanEqualConstraint</code> , which requires that a boolean axis equal either true or false (depending on user input).
<code>EnumConstraint</code>	A superclass in which constraints on enum values must extend. Currently, the only subclass is <code>EnumEqualConstraint</code> , which requires that an enum axis be equal to a user-specified value.
<code>IntegerConstraint</code>	A superclass in which constraints on integer values must extend. Subclasses include <code>IntegerLessThanConstraint</code> , <code>IntegerGreaterThanConstraint</code> , and <code>IntegerEqualConstraint</code> , which map to the obvious comparators for integers.
<code>NullConstraint</code>	This class signifies that an axis cannot be constrained. This makes the most sense for axes that provide an average value - placing a constraint on an average doesn't make sense because you are not able to include or discard a particular operation based on what its effects would be on the average of all operations.
<code>PathConstraint</code>	A superclass in which constraints on file path values must extend. Currently, the only subclass is <code>PathDescendantConstraint</code> , which requires that a path value must be a descendant of the specified path (it must be contained within it). This only applies to paths on the Delphix Engine itself, and all paths used must be canonical Unix paths starting from the root of the filesystem.
<code>StringConstraint</code>	A superclass in which constraints on string values must extend. Currently, the only subclass is <code>StringEqualsConstraint</code> , which requires that a string value must equal a user-specified string.

Statistic Slice Commands

Command	Description and Usage Examples
<code>getData</code>	This is used to fetch data from a statistic slice which has been collecting data for a while. It returns a datapoint set, which is composed of datapoint streams, which contain datapoints. For a full description, see the Performance Analytics Tool Overview .

Command	Description and Usage Examples
<code>rememberRange</code>	This is used to ensure that data collected during an ongoing investigation doesn't get deleted unexpectedly. If this is not used, data is only guaranteed to be persisted for 24 hours. If it is used, data will be remembered until a corresponding call to <code>stopRememberingRange</code> is made.
<code>stopRememberingRange</code>	This is used to allow previously-remembered data to be forgotten. The data will be forgotten on the same schedule as brand new data, so you will have at least 24 hours before data which you have stopped remembering is deleted. This undoes the <code>rememberRange</code> operation.
<code>pause</code>	This command pauses the collection of a statistic slice, causing no data to be collected until <code>resume</code> is called.
<code>resume</code>	This command resumes the collection of a statistic slice, undoing a <code>pause</code> operation.

Related Topic

- [Performance Analytics Tool Overview](#)
- [Performance Analytics Statistics Reference](#)

Performance Analytics Case Study: Using a Single Statistic

Overview

The Delphix Engine uses the Network File System (NFS) as the transport for Oracle installations. An increase in the NFS latency could be causing sluggishness in your applications running on top of Virtual Databases. This case study illustrates how this pathology can be root caused using the analytics infrastructure. This performance investigation uses one statistic to debug the issue, and utilizes the many axes of that statistic to filter down the probable cause of the issue. This technique uses an approach of iteratively drilling down by inspecting new axes of a single statistic and filtering the data to only include information about the operations that appear slow. This technique is valuable for determining which use patterns of a resource might be causing the system to be sluggish. If you isolate a performance issue using this approach but aren't sure what is causing it or how to fix it, Delphix Support can provide assistance for your investigation.

The following example inspects the statistic which provides information about NFS I/O operations on the Delphix Engine. This statistic can be collected a maximum of once every second, and the axes it can collect, among others, are:

- **latency**, a histogram of wait times between NFS requests and NFS responses
- **size**, a histogram of the NFS I/O sizes requested
- **op**, whether the NFS requests were reads or writes
- **client**, the network address of the NFS client which was making requests

 Roughly the same performance information can be obtained from the iSCSI interface as well.

Investigation

Because the NFS layer sits above the disk layer, all NFS operations that use the disk synchronously (synchronous writes and uncached reads) will have latencies that are slightly higher than those of their corresponding disk operations. Usually, because disks have very high seek times compared to the time the NFS server spends on CPU, disk operations are responsible for almost all of the latency of these NFS operations. In the graphical representation, you can see this by looking at how the slower cluster of NFS latencies (around 2ms-8ms) have similar latencies to the median of the disk I/O (around 2ms-4ms). Another discrepancy between the two plots is that the number of disk operations is much lower than the corresponding number of NFS operations. This is because the Delphix filesystem batches together write operations to improve performance.

If database performance is not satisfactory and almost all of the NFS operation time is spent waiting for the disks, it suggests that the disk is the slowest piece of the I/O stack. In this case, disk resources (the number of IOPS to the disks, the free space on the disks, and the disk throughput) should be investigated more thoroughly to determine if adding more capacity or a faster disk would improve performance. However, care must be taken when arriving at these conclusions, as a shortage of memory or a recently-rebooted machine can also cause the disk to be used more heavily due to fewer cache hits.

Sometimes, disk operations will not make up all of the latency, which suggests that something between the NFS server and the disk (namely, something in the Delphix Engine) is taking a long time to complete its work. If this is the case, it is valuable to check whether the Delphix Engine is resource-constrained, and the most common areas of constraint internal to the Delphix Engine are CPU and memory. If either of those is too limited, you should investigate whether expanding the resource would improve performance. If no resources appear to be constrained or more investigation is necessary to convince you that adding resources would help the issue, Delphix Support is available to help debug these issues.

While using this technique, you should take care to recognize the limitations that caching places on how performance data can be interpreted. In this example, the Delphix Engine uses a caching layer for the data it stores, so asynchronous NFS writes will not go to disk quickly because they are being queued into larger batches, and

cached NFS reads won't use the disk at all. This causes these types of NFS operations to return much more quickly than any disk operations are able to, resulting in a very large number of low-latency NFS operations in the graph above. For this reason, caching typically creates a bimodal distribution in the NFS latency histograms, where the first cluster of latencies is associated with operations that only hit the cache, and the second cluster of latencies is associated with fully or partially uncached operations. In this case, cached NFS operations should not be compared to the disk latencies because they are unrelated. It is possible to use techniques described in the first example to filter out some of the unrelated operations to allow a more accurate mapping between disk and NFS latencies.

1. Begin the performance investigation by examining some high-level statistics such as **latency**.
 - a. Create a slice with statistic type **NFS_OPS**.
 - b. Set the slice to collect the **latency** axis.
 - c. Do not add any constraints.
 - d. Set the collection interval.
Anything over one second will work, but ten seconds gives good data resolution and will not use a lot of storage to persist the data that is collected. The rest of this example will assume a collection period of ten seconds for all other slices, but any value could be used.

```
/analytics

create
set name=step1
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=latency
commit
```

This will collect a time-series of histograms describing NFS latencies as measured from inside the Delphix Engine, where each histogram shows how many NFS I/O operations fell into each latency bucket during every ten-second interval. After a short period of time, read the data from the statistic slice:

```
select step1
getData
setopt format=json
commit
setopt format=text
```

The `setopt` steps are optional but allow you to see the output better via the CLI. The output looks like this:

```
{
  "type": "DatapointSet",
  "collectionEvents": [],
  "datapointStreams": [{
    "type": "NfsOpsDatapointStream",
    "datapoints": [{
      "type": "IoOpsDatapoint",
      "latency": {
        "32768": "16",
```

```

        "65536": "10"
      },
      "timestamp": "2013-05-14T15:51:40.000Z"
    }, ...]
  ]],
  "resolution": 10
}

```

The data is returned as a set of datapoint streams. Streams hold the fields which are shared by all the datapoints they contain. Later on, in this example, the **opt** and **client** fields will be added to the streams, and multiple streams will be returned. Streams are described in more detail in [Performance Analytics Tool Overview](#). The `resolution` field indicates the number of seconds that corresponds to each datapoint, which in our case matches the requested `collectionInterval`. The `collectionEvents` field is not used in this example, but lists when the slice was paused and resumed, to distinguish between moments when no data was collected because the slice was paused, and moments when there was no data to collect.

2. If the latency distributions show some slow NFS operations, the next step would be to determine whether the slow operations are reads or writes.
 - a. Specify a new **NFS_OPS** slice to collect this by collecting the **op** and **latency** axes.
 - b. To limit output to the long-running operations, create a constraint on the **latency** axis that prohibits the collection of data on operations with latency less than 100ms.

```

/analytics

create
set name=step2
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=op,latency

edit axisConstraints.0
set axisName=latency
set type=IntegerGreaterThanConstraint
set greaterThan=100000000
back

commit

```

The `greaterThan` field is 100ms converted into nanoseconds.

Reading the data proceeds in the same way as the first step, but there will be two streams of datapoints, one where `op=write`, and one where `op=read`.

 Because we constrained output to operations with latencies higher than 100ms, none of the latency histograms will all have any buckets for latencies lower than 100ms.

3. After inspecting the two data streams, you might find that almost all slow operations are writes, so it could be valuable to determine which clients are requesting the slow writes, and how large each of the writes is.
 - a. To collect this data, create a new **NFS_OPS** slice which collects the **size** and **client** axes.

- b. Add constraints ensuring that the **op** axis should be constrained to only collect data for **write** operations, and the **latency** axis should be constrained to filter operations taking less than 100ms.

i Because the constraint on the **op** axis dictates that it will always have the value **write**, it is not necessary to collect the **op** axis anymore.

```

/analytics

create
set name=step3
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=size,client

edit axisConstraints.0
set type=IntegerGreaterThanConstraint
set axisName=latency
set greaterThan=100000000
back

edit EnumEqualConstraint
set type=StringEqualConstraint
set axisName=op
set equals=write
back

commit

```

Reading the data proceeds in the same way as the first two steps, but there will be one stream for every NFS client. The dataset collected by this will consist of a set of streams, one corresponding to each NFS client, and each stream will be a time-series of histograms showing write sizes that occurred during each ten-second interval.

Continuing to use this approach will allow you to narrow down the slow writes to a particular NFS client, and you may be able to tune that client in some way to speed it up.

Related Topics

- [Performance Analytics Tool Overview](#)
- [Performance Analytics Case Study: Using Multiple Statistics](#)

Performance Analytics Case Study: Using Multiple Statistics

Overview

This case study illustrates an investigation involving more than one metric. In typical performance investigations, you will need to peel out multiple layers of the stack in order to observe the component causing the actual performance pathology. This case study specifically examines sluggish application performance caused due to slow IO responses from the disk subsystem. This example will demonstrate a technique of looking at the performance of each layer in the I/O stack to find which layer is responsible for the most latency, then looking for constrained resources that the layer might need to access. This technique is valuable for finding the most-constrained resource in the system, potentially giving actionable information about resources that can be expanded to increase performance.

For the following example, we will inspect latency at two layers: the Network File System (NFS) layer on the Delphix Engine, and the disk layer below it. Both of these layers provide the **latency** axis, which gives a histogram of wait times for the clients of each layer.

Investigation

The analytics infrastructure enables users to observe the latency of multiple layers of the software stack. This investigation will examine the latency of both layers, and then draw conclusions about the differences between the two.

Setup

To measure this data, create two slices. When attempting to correlate data between two different statistics, it can be easier to determine causation when collecting data at a relatively high frequency. The fastest that each of these statistics will collect data is once per second, so that is the value used.

1. A slice collecting the **latency** axis for the statistic type **NFS_OPS**.

```
/analytics  
  
create  
set name=slice1  
set statisticType=NFS_OPS  
set collectionInterval=1  
set collectionAxes=latency  
commit
```

2. A slice collecting the **latency** axis for the statistic type **DISK_OPS**.

```
/analytics  
  
create  
set name=slice2  
set statisticType=DISK_OPS  
set collectionInterval=1  
set collectionAxes=latency  
commit
```

After a short period of time, read the data from the first statistic slice.

```
select slice2
getData
setopt format=json
commit
setopt format=text
```

The same process works for the second slice. The `setopt` steps allow you to see the output better via the CLI. The output for the first slice might look like this:

```
{
  "type": "DatapointSet",
  "collectionEvents": [],
  "datapointStreams": [{
    "type": "NfsOpsDatapointStream",
    "datapoints": [{
      "type": "IoOpsDatapoint",
      "latency": {
        "512": "100",
        "1024": "308",
        "2048": "901",
        "4096": "10159",
        "8192": "2720",
        "16384": "642",
        "32768": "270",
        "65536": "50",
        "131072": "11",
        "524288": "64",
        "1048576": "102",
        "2097152": "197",
        "4194304": "415",
        "8388608": "320",
        "16777216": "50",
        "33554432": "20",
        "67108864": "9",
        "268435456": "2"
      },
      "timestamp": "2013-05-14T15:51:40.000Z"
    }, {
      "type": "IoOpsDatapoint",
      "latency": {
        "512": "55",
        "1024": "130",
        "2048": "720",
        "4096": "6500",
        "8192": "1598",
        "16384": "331",
        "32768": "327",
        "65536": "40",
```

```

        "131072": "14",
        "262144": "87",
        "524288": "42",
        "1048576": "97",
        "2097152": "662",
        "4194304": "345",
        "8388608": "280",
        "16777216": "22",
        "33554432": "15",
        "134217728": "1"
    },
    "timestamp": "2013-05-14T15:51:41.000Z"
  }, ...]
}],
"resolution": 1
}

```

For the second slice, it might look like this:

```

{
  "type": "DatapointSet",
  "collectionEvents": [],
  "datapointStreams": [{
    "type": "DiskOpsDatapointStream",
    "datapoints": [{
      "type": "IoOpsDatapoint",
      "latency": {
        "262144": "1",
        "524288": "11",
        "1048576": "13",
        "2097152": "34",
        "4194304": "7",
        "8388608": "2",
        "16777216": "3",
        "33554432": "1"
      }
    },
    "timestamp": "2013-05-14T15:51:40.000Z"
  }, {
    "type": "IoOpsDatapoint",
    "latency": {
      "262144": "5",
      "524288": "10",
      "1048576": "14",
      "2097152": "26",
      "4194304": "7",
      "8388608": "4",
      "16777216": "2"
    },
    "timestamp": "2013-05-14T15:51:41.000Z"
  }, ...]
}],
}

```



```

    "resolution": 1
  }

```

The data is returned as a set of datapoint streams. Streams hold the fields that would otherwise be shared by all the datapoints they contain, but only one is used in this example because there are no such fields. Streams are discussed in more detail in the [Performance Analytics Tool Overview](#). The `resolution` field indicates how many seconds each datapoint corresponds to, which in our case matches the requested `collectionInterval`. The `collectionEvents` field is not used in this example, but lists when the slice was paused and resumed to distinguish between moments when no data was collected because the slice was paused, and moments when there was no data to collect.

Graphically, these four histograms across two seconds look like this:



Analysis

Because the NFS layer sits above the disk layer, all NFS operations that use the disk synchronously (synchronous writes and uncached reads) will have latencies that are slightly higher than those of their corresponding disk operations. Usually, because disks have very high seek times compared to the time the NFS server spends on CPU, disk operations are responsible for almost all of the latency of these NFS operations. In the graphical representation, you can see this by looking at how the slower cluster of NFS latencies (around 2ms-8ms) have similar latencies to the median of the disk I/O (around 2ms-4ms). Another discrepancy between the two plots is that the number of disk operations is much lower than the corresponding number of NFS operations. This is because the Delphix filesystem batches together write operations to improve performance.

If database performance is not satisfactory and almost all of the NFS operation time is spent waiting for the disks, it suggests that the disk is the slowest piece of the I/O stack. In this case, disk resources (the number of IOPS to the disks, the free space on the disks, and the disk throughput) should be investigated more thoroughly to determine if adding more capacity or a faster disk would improve performance. However, care must be taken when arriving at these conclusions, as a shortage of memory or a recently-rebooted machine can also cause the disk to be used more heavily due to fewer cache hits.

Sometimes, disk operations will not make up all of the latency, which suggests that something between the NFS server and the disk (namely, something in the Delphix Engine) is taking a long time to complete its work. If this is the case, it is valuable to check whether the Delphix Engine is resource-constrained, and the most common areas of constraint internal to the Delphix Engine are CPU and memory. If either of those is too limited, you should investigate whether expanding the resource would improve performance. If no resources appear to be constrained or more investigation is necessary to convince you that adding resources would help the issue, Delphix Support is available to help debug these issues.

While using this technique, you should take care to recognize the limitations that caching places on how performance data can be interpreted. In this example, the Delphix Engine uses a caching layer for the data it stores, so asynchronous NFS writes will not go to disk quickly because they are being queued into larger batches, and cached NFS reads won't use the disk at all. This causes these types of NFS operations to return much more quickly than any disk operations are able to, resulting in a very large number of low-latency NFS operations in the graph above. For this reason, caching typically creates a bimodal distribution in the NFS latency histograms, where the first cluster of latencies is associated with operations that only hit the cache, and the second cluster of latencies is associated with fully or partially uncached operations. In this case, cached NFS operations should not be compared to the disk latencies because they are unrelated. It is possible to use techniques described in the first example to filter out some of the unrelated operations to allow a more accurate mapping between disk and NFS latencies.

1. Begin the performance investigation by examining some high-level statistics such as **latency**.
 - a. Create a slice with statistic type **NFS_OPS**.
 - b. Set the slice to collect the **latency** axis.
 - c. Do not add any constraints.
 - d. Set the collection interval.
Anything over one second will work, but ten seconds gives good data resolution and will not use a lot of storage to persist the data that is collected. The rest of this example will assume a collection period of ten seconds for all other slices, but any value could be used.

```
/analytics

create
set name=step1
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=latency
```

```
commit
```

This will collect a time series of histograms describing NFS latencies as measured from inside the Delphix Engine, where each histogram shows how many NFS I/O operations fell into each latency bucket during every ten-second interval. After a short period of time, read the data from the statistic slice:

```
select step1
getData
setopt format=json
commit
setopt format=text
```

The `setopt` steps are optional but allow you to see the output better via the CLI. The output looks like this:

```
{
  "type": "DatapointSet",
  "collectionEvents": [],
  "datapointStreams": [{
    "type": "NfsOpsDatapointStream",
    "datapoints": [{
      "type": "IoOpsDatapoint",
      "latency": {
        "32768": "16",
        "65536": "10"
      },
      "timestamp": "2013-05-14T15:51:40.000Z"
    }, ...]
  }],
  "resolution": 10
}
```

The data is returned as a set of datapoint streams. Streams hold the fields which are shared by all the data points they contain. Later on, in this example, the **opt** and **client** fields will be added to the streams, and multiple streams will be returned. Streams are described in more detail in [Performance Analytics Tool Overview](#). The `resolution` field indicates the number of seconds that corresponds to each datapoint, which in our case matches the requested `collectionInterval`. The `collectionEvents` field is not used in this example, but lists when the slice was paused and resumed, to distinguish between moments when no data was collected because the slice was paused, and moments when there was no data to collect.

2. If the latency distributions show some slow NFS operations, the next step would be to determine whether the slow operations are reads or write.
 - a. Specify a new **NFS_OPS** slice to collect this by collecting the **op** and **latency** axes.
 - b. To limit output to the long-running operations, create a constraint on the **latency** axis that prohibits the collection of data on operations with latency less than 100ms.

```

/analytics

create
set name=step2
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=op,latency

edit axisConstraints.0
set axisName=latency
set type=IntegerGreaterThanConstraint
set greaterThan=100000000
back

commit

```

The `greaterThan` field is 100ms converted into nanoseconds.

Reading the data proceeds in the same way as the first step, but there will be two streams of data points, one where `op=write`, and one where `op=read`.

i Because we constrained output to operations with latencies higher than 100ms, none of the latency histograms will all have any buckets for latencies lower than 100ms.

3. After inspecting the two data streams, you might find that almost all slow operations are writes, so it could be valuable to determine which clients are requesting the slow writes, and how large each of the writes is.
 - a. To collect this data, create a new **NFS_OPS** slice which collects the **size** and **client** axes.
 - b. Add constraints ensuring that the **op** axis should be constrained to only collect data for **write** operations, and the **latency** axis should be constrained to filter operations taking less than 100ms.

i Because the constraint on the **op** axis dictates that it will always have the value **write**, it is not necessary to collect the **op** axis anymore.

```

/analytics

create
set name=step3
set statisticType=NFS_OPS
set collectionInterval=10
set collectionAxes=size,client

edit axisConstraints.0
set type=IntegerGreaterThanConstraint
set axisName=latency
set greaterThan=100000000
back

edit EnumEqualConstraint

```

```
set type=StringEqualConstraint
set axisName=op
set equals=write
back

commit
```

Reading the data proceeds in the same way as the first two steps, but there will be one stream for every NFS client. The dataset collected by this will consist of a set of streams, one corresponding to each NFS client, and each stream will be a time series of histograms showing write sizes that occurred during each ten-second interval.

Continuing to use this approach will allow you to narrow down the slow writes to a particular NFS client, and you may be able to tune that client in some way to speed it up.

Related Topics

- [Performance Analytics Tool Overview](#)
- [Performance Analytics Case Study: Using a Single Statistic](#)

Storage Performance Configuration Options

This section covers the following topics:

- [Optimal Storage Configuration Parameters for the Delphix Engine](#)
- [Storage Performance Test Tool \(fio\)](#)
- [Storage Performance Expectations and Troubleshooting](#)
- [Storage Performance Test Tool notes - Restricted](#)

Optimal Storage Configuration Parameters for the Delphix Engine

This topic describes minimum capacity and throughput requirements for storage devices used with the Delphix Engine.

Storage for the Delphix Engine must be able to sustain the aggregated Input/Output Operations Per Second (IOPS) and throughput (MBPS) requirements of all its Virtual Databases. Throughput required for data source synchronization (SnapSync and LogSync) must also be supported.

The Delphix Engine requires storage for:

Item	Description
A copy of each Source Database	The copies are compressed.
Unique Block Changes in VDBs	When changes are made to a VDB, the Delphix Engine stores the changes in new blocks associated with only that VDB. The new blocks are compressed.
Timeflow for dSources and VDBs	The TimeFlow kept for each dSource and VDB comprises snapshots of the database (blocks changed since the previous snapshot) and archive logs. The retention period for this history of changes is determined by policies established for each dSource and VDB. The TimeFlow is compressed.

In addition to the storage for these items, the Delphix Engine requires 30% free space in its storage for the best performance. See [An Overview of Capacity and Performance Information](#) and related topics for more details on managing capacity for the Delphix Engine.

Best practices for storage performance include:

- Initial storage is equal to the size of the physical source databases. For high redo rates and/or high DB change rates, allocate an additional 10-20% storage.
- Add storage when storage capacity approaches 30% free
- Use physical LUNS allocated from storage pools or RAID groups that are configured for availability
- Never share physical LUNS between the Delphix Engine and other storage clients.
- Keep all physical LUNS the same size.
- Provision storage using VMDKs or RDMs operating in virtual compatibility mode.
- VMDKs should be **Thick Provisioned, Lazy Zeroed**. The underlying physical LUNS can be thin provisioned.
- Physical LUNS used for RDMs should be thick provisioned.
- Measure or estimate the required IOPS and manage the storage disks to provide this capacity. It is common to use larger numbers of spindles to provide the IOPS required.
- Physical LUNS carved from RAID 1+0 groups or pools with dedicated spindles provide higher IOPS performance than other configurations
- Maximize Delphix Engine vRAM for a larger system cache to service reads

Example

There are two production dSources, totaling 5 TB in size. 5 VDBs will be created for each. The Sum of read and write rates on the production source database is moderate (1000 IOPS), the sum of VDB read rate is moderate (950 IOPS), and the VDB update rate is low (50 IOPS).

- Initial storage equal to 5TB, provisioned as 5 x 1 TB physical LUNs, Thin Provisioned. Allow for expansion of the LUNs to 2TB.
- Provision as 5 x 950 GB Virtual Disks. VMDKs must be Thick Provisioned, Lazy Zeroed. Using 1 TB LUNs allows expansion to 2 TB (ESX 5.1 limit).
- The storage provisioned to the Delphix Engine storage must be able to sustain 1000 IOPs (950 + 50). For this reason, each physical LUN provisioned to the Delphix Engine must be capable of sustaining 200 IOPs. IOPs on the source databases are not relevant to the Delphix Engine.
- 64GB Delphix Engine vRAM for a large system cache

Related Topics

- [Storage Performance Expectations and Troubleshooting](#)
- [An Overview of Capacity and Performance Information](#)

Storage Performance Test Tool (fio)

Overview

This fio-based Storage Performance Tool executes a synthetic workload to evaluate the performance characteristics of the storage assigned to the Delphix Engine. The Storage Performance Tool is a feature that is only available from the command-line

Prerequisites

Prior to setting up the Delphix Engine, the admin can login to the Delphix CLI using a **sysadmin** account to launch the Storage Performance Tool. Because the test is destructive, it will only run against storage which has not been allocated to Delphix for use by the engine.

Running the Storage Test via CLI

1. Login as the sysadmin user to the Delphix Engine CLI.
 - a. If the Delphix Engine has not been setup yet, the **network setup** prompt appears. Discard the command.

```
delphix network setup update *> discard delphix>
```

2. Create a storage test.

```
delphix> storage test
delphix storage test> create
delphix storage test create *>
```

3. Use 'get' to see other optional arguments. Modify the test parameters as needed and **commit** to start the test.

```
delphix storage test create *> get
type: StorageTestParameters
devices: (unset)
duration: 120
initializeDevices: true
initializeEntireDevice: false
testRegion: 128GB
tests: ALL
delphix storage test create *> commit
STORAGE_TEST-1
Dispatched job JOB-1
STORAGE_TEST_EXECUTE job started for "SYSTEM".
Initializing storage test.
Starting storage device initialization.
ETA: 1:28:44.
Storage device initialization complete.
Starting storage benchmarking.
Starting random read workload with 4 KB block size and 8 jobs.
Starting random read workload with 4 KB block size and 16 jobs.
```

```

Starting random read workload with 4 KB block size and 32 jobs.
Starting random read workload with 4 KB block size and 64 jobs.
Starting random read workload with 8 KB block size and 8 jobs.
Starting random read workload with 8 KB block size and 16 jobs.
Starting random read workload with 8 KB block size and 32 jobs.
Starting random read workload with 8 KB block size and 64 jobs.
Starting sequential write workload with 1 KB block size and 4 jobs.
Starting sequential write workload with 4 KB block size and 4 jobs.
Starting sequential write workload with 8 KB block size and 4 jobs.
Starting sequential write workload with 16 KB block size and 4 jobs.
Starting sequential write workload with 32 KB block size and 4 jobs.
Starting sequential write workload with 64 KB block size and 4 jobs.
Starting sequential write workload with 128 KB block size and 4 jobs.
Starting sequential write workload with 1024 KB block size and 4 jobs.
Starting sequential write workload with 1 KB block size and 16 jobs.
Starting sequential write workload with 4 KB block size and 16 jobs.
Starting sequential write workload with 8 KB block size and 16 jobs.
Starting sequential write workload with 16 KB block size and 16 jobs.
Starting sequential write workload with 32 KB block size and 16 jobs.
Starting sequential write workload with 64 KB block size and 16 jobs.
Starting sequential write workload with 128 KB block size and 16 jobs.
Starting sequential write workload with 1024 KB block size and 16 jobs.
Starting sequential read workload with 64 KB block size and 4 jobs.
Starting sequential read workload with 64 KB block size and 8 jobs.
Starting sequential read workload with 64 KB block size and 16 jobs.
Starting sequential read workload with 64 KB block size and 32 jobs.
Starting sequential read workload with 64 KB block size and 64 jobs.
Starting sequential read workload with 128 KB block size and 4 jobs.
Starting sequential read workload with 128 KB block size and 8 jobs.
Starting sequential read workload with 128 KB block size and 16 jobs.
Starting sequential read workload with 128 KB block size and 32 jobs.
Starting sequential read workload with 128 KB block size and 64 jobs.
Starting sequential read workload with 1024 KB block size and 4 jobs.
Starting sequential read workload with 1024 KB block size and 8 jobs.
Starting sequential read workload with 1024 KB block size and 16 jobs.
Starting sequential read workload with 1024 KB block size and 32 jobs.
Starting sequential read workload with 1024 KB block size and 64 jobs.
Storage benchmarking complete.
Generating results.
Storage test completed successfully.
STORAGE_TEST_EXECUTE job for "SYSTEM" completed successfully.
delphix storage test>

```

4. The job will be submitted and visible in the Delphix Management application.
5. Retrieve the test results

```

delphix storage test> select STORAGE_TEST-1
delphix storage test 'STORAGE_TEST-1'> result
delphix storage test 'STORAGE_TEST-1' result *> commit
Test Results
-----
Test ID:          1

```

Test System UUID: 564dc710-7bb1-c064-12c2-2659032acf1b
 Start Time: 03-Feb-2015 10:52:31 -0800
 End Time: 03-Feb-2015 12:20:25 -0800

Test Grades:

Test Name	Latency			Load Scaling	
	Average	95th %ile	Grade	Scaling	Grade
Random 8K Reads w/ 16 jobs	2.16	4.77	A-	0.89	poor
Random 4K Reads w/ 16 jobs	1.62	3.73	A	0.54	fair
Sequential 1M Reads w/ 4 jobs	62.60	182.00	D	1.40	bad
Sequential 1K Writes w/ 4 jobs	1.30	2.61	C	0.07	good
Sequential 128K Writes w/ 4 jobs	10.19	26.00	D	1.35	bad

Grading Key:

Test Name	Grade:	A+	A	A-	B	B-	C	C-	D
Small Random Reads		2.0	4.0	6.0	8.0	10.0	12.0	14.0	> 14.0
Large Seq Reads		12.0	14.0	16.0	18.0	20.0	22.0	24.0	> 24.0
Small Seq Writes		0.5	1.0	1.5	2.0	2.5	3.0	3.5	> 3.5
Large Seq Writes		2.0	4.0	6.0	8.0	10.0	12.0	14.0	> 14.0

IO Summary:

Test Name	IOPS	Throughput (MBps)	Latency (msec)	
			Average	Min
Max	StdDev			
Rand 4K Reads w/ 8 Jobs	15703	61.34	0.50	0.05
754.74	1.72			
Rand 4K Reads w/ 16 Jobs	15631	61.06	1.00	0.05
1347.10	5.12			
Rand 4K Reads w/ 32 Jobs	15972	62.39	1.95	0.05
1231.40	17.56			
Rand 4K Reads w/ 64 Jobs	17341	67.74	3.62	0.05
1750.10	30.09			
Rand 8K Reads w/ 8 Jobs	15151	118.37	0.52	0.05
45.18	0.27			
Rand 8K Reads w/ 16 Jobs	16457	128.58	0.95	0.05
501.90	3.57			
Rand 8K Reads w/ 32 Jobs	16908	132.10	1.84	0.05
1336.10	16.93			
Rand 8K Reads w/ 64 Jobs	16865	131.76	3.71	0.05
1505.50	30.03			

Seq 1K Writes w/ 4 Jobs	22053	21.54	0.18	0.04
168.14 0.26				
Seq 4K Writes w/ 4 Jobs	24937	97.41	0.16	0.04
152.17 0.27				
Seq 8K Writes w/ 4 Jobs	22946	179.27	0.17	0.04
120.19 0.28				
Seq 16K Writes w/ 4 Jobs	18003	281.31	0.22	0.05
81.24 0.26				
Seq 32K Writes w/ 4 Jobs	12993	406.05	0.30	0.05
40.33 0.38				
Seq 64K Writes w/ 4 Jobs	6429	401.83	0.62	0.06
116.19 2.26				
Seq 128K Writes w/ 4 Jobs	3614	451.86	1.10	0.08
200.12 4.75				
Seq 1M Writes w/ 4 Jobs	388	388.83	10.28	0.27
832.57 42.24				
Seq 1K Writes w/ 16 Jobs	25965	25.36	0.60	0.04
814.84 6.86				
Seq 4K Writes w/ 16 Jobs	25610	100.04	0.61	0.04
1022.50 7.76				
Seq 8K Writes w/ 16 Jobs	25183	196.75	0.62	0.04
910.55 7.91				
Seq 16K Writes w/ 16 Jobs	23433	366.14	0.66	0.04
948.57 8.05				
Seq 32K Writes w/ 16 Jobs	19327	604.00	0.81	0.05
1180.50 8.10				
Seq 64K Writes w/ 16 Jobs	9313	582.08	1.71	0.06
711.96 4.40				
Seq 128K Writes w/ 16 Jobs	3369	421.14	4.75	0.08
69.12 4.32				
Seq 1M Writes w/ 16 Jobs	481	481.06	33.22	0.27
269.88 32.05				
Seq 64K Reads w/ 4 Jobs	16912	1057.20	0.23	0.05
40.36 0.15				
Seq 64K Reads w/ 8 Jobs	18862	1178.10	0.42	0.05
78.57 0.41				
Seq 64K Reads w/ 16 Jobs	20352	1272.50	0.77	0.06
900.81 7.41				
Seq 64K Reads w/ 32 Jobs	20750	1296.10	1.50	0.06
1231.60 20.02				
Seq 64K Reads w/ 64 Jobs	21146	1321.70	2.95	0.06
2440.30 34.37				
Seq 128K Reads w/ 4 Jobs	11649	1456.30	0.34	0.06
53.66 0.25				
Seq 128K Reads w/ 8 Jobs	15995	1999.50	0.49	0.06
32.21 0.42				
Seq 128K Reads w/ 16 Jobs	17413	2176.80	0.90	0.07
1057.60 6.46				
Seq 128K Reads w/ 32 Jobs	17874	2234.30	1.76	0.07
1355.40 19.87				
Seq 128K Reads w/ 64 Jobs	17523	2190.50	3.58	0.07
1926.20 36.79				

Seq 1M Reads w/ 4 Jobs	1404	1404.20	2.84	0.31
64.38 0.75				
Seq 1M Reads w/ 8 Jobs	2360	2360.70	3.38	0.32
17.60 0.46				
Seq 1M Reads w/ 16 Jobs	3876	3876.50	4.10	0.33
429.44 3.20				
Seq 1M Reads w/ 32 Jobs	4732	4732.60	6.69	0.29
1305.70 34.64				
Seq 1M Reads w/ 64 Jobs	4730	4730.10	13.33	0.32
1847.90 54.39				

IO Histogram:

Test Name				us50	us100	us250	us500	ms1	ms2	ms4
ms10	ms20	ms50	ms100	ms250	ms500	s1	s2	s5		
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----	-----
Rand 4K Reads w/ 8 Jobs				0	0	0	2	46	41	7
3 1 0 0				0	0	0	0	0		
Rand 4K Reads w/ 16 Jobs				0	0	0	0	39	47	10
3 1 0 0				0	0	0	0	0		
Rand 4K Reads w/ 32 Jobs				0	0	0	0	6	64	22
6 2 0 0				0	0	0	0	0		
Rand 4K Reads w/ 64 Jobs				0	0	0	0	0	4	75
16 3 1 0				0	0	0	0	0		
Rand 8K Reads w/ 8 Jobs				0	0	0	0	41	49	7
2 1 0 0				0	0	0	0	0		
Rand 8K Reads w/ 16 Jobs				0	0	0	0	8	66	20
4 2 1 0				0	0	0	0	0		
Rand 8K Reads w/ 32 Jobs				0	0	0	0	0	5	72
18 3 2 0				0	0	0	0	0		
Rand 8K Reads w/ 64 Jobs				0	0	0	0	0	0	3
85 8 4 1				0	0	0	0	0		
Seq 1K Writes w/ 4 Jobs				0	0	0	4	53	36	4
2 1 0 0				0	0	0	0	0		
Seq 4K Writes w/ 4 Jobs				0	0	0	2	44	44	6
3 1 0 0				0	0	0	0	0		
Seq 8K Writes w/ 4 Jobs				0	0	0	1	41	47	7
3 1 0 0				0	0	0	0	0		
Seq 16K Writes w/ 4 Jobs				0	0	0	0	27	57	10
3 1 0 0				0	0	0	0	0		
Seq 32K Writes w/ 4 Jobs				0	0	0	0	4	55	30
8 2 1 0				0	0	0	0	0		
Seq 64K Writes w/ 4 Jobs				0	0	0	0	0	1	56
33 7 3 0				0	0	0	0	0		
Seq 128K Writes w/ 4 Jobs				0	0	0	0	0	0	0
76 16 5 2				0	0	0	0	0		
Seq 1M Writes w/ 4 Jobs				0	0	0	0	0	0	0
0 0 24 57				14	4	0	0	0		

```

Seq 1K Writes w/ 16 Jobs      0      0      0      1      55      34      6
2      1      0      0      0      0      0      0      0
Seq 4K Writes w/ 16 Jobs      0      0      0      1      46      42      7
3      1      0      0      0      0      0      0      0
Seq 8K Writes w/ 16 Jobs      0      0      0      0      3      43      38
12     2      1      0      0      0      0      0      0
Seq 16K Writes w/ 16 Jobs     0      0      0      0      0      4      67
23     4      2      0      0      0      0      0      0
Seq 32K Writes w/ 16 Jobs     0      0      0      0      0      0      1
74     16     8      2      0      0      0      0      0
Seq 64K Writes w/ 16 Jobs     0      0      0      0      0      0      0
2      81     12     3      1      0      0      0      0
Seq 128K Writes w/ 16 Jobs    0      0      0      0      0      0      0
1      2      85     7      4      0      0      0      0
Seq 1M Writes w/ 16 Jobs      0      0      0      0      0      0      0
0      0      1      4     45     38     9      3      0
Seq 64K Reads w/ 4 Jobs       0      0      0      0      0      29     59
9      2      1      0      0      0      0      0      0
Seq 64K Reads w/ 8 Jobs       0      0      0      0      0      0      15
74     8      3      1      0      0      0      0      0
Seq 64K Reads w/ 16 Jobs      0      0      0      0      0      0      0
14     53     27     5      1      0      0      0      0
Seq 64K Reads w/ 32 Jobs      0      0      0      0      0      0      0
1      27     59     8      4      0      0      0      0
Seq 64K Reads w/ 64 Jobs      0      0      0      0      0      0      0
0      1      29     42     25     2      0      0      0
Seq 128K Reads w/ 4 Jobs      0      0      0      0      0      0      10
75     9      5      1      0      0      0      0      0
Seq 128K Reads w/ 8 Jobs      0      0      0      0      0      0      0
64     29     5      2      0      0      0      0      0
Seq 128K Reads w/ 16 Jobs     0      0      0      0      0      0      0
1      45     45     6      2      0      0      0      0
Seq 128K Reads w/ 32 Jobs     0      0      0      0      0      0      0
0      1      65     24     8      1      0      0      0
Seq 128K Reads w/ 64 Jobs     0      0      0      0      0      0      0
0      1      8      54     29     5      1      0      0
Seq 1M Reads w/ 4 Jobs        0      0      0      0      0      0      0
0      0      66     23     8      2      0      0      0
Seq 1M Reads w/ 8 Jobs        0      0      0      0      0      0      0
0      0      1      33     52     11     3      0      0
Seq 1M Reads w/ 16 Jobs       0      0      0      0      0      0      0
0      0      1      5      70     15     8      1      0
Seq 1M Reads w/ 32 Jobs       0      0      0      0      0      0      0
0      0      1      4      19     58     11     6      1
Seq 1M Reads w/ 64 Jobs       0      0      0      0      0      0      0
0      0      1      2      10     40     32     11     2

```

```
delphix storage test 'STORAGE_TEST-1'>
```

Related Topics

- [Optimal Storage Configuration Parameters for the Delphix Engine](#)
- [Storage Performance Expectations and Troubleshooting](#)

Storage Performance Expectations and Troubleshooting

Overview

When you initially set up a Delphix Engine, there is a one-time opportunity (and requirement) to run your storage performance tests on unconfigured storage. Because of the test method we use, we cannot run it again in the same way after the engine is set up. When run from a Windows target host, [DiskSpd](#) can help establish a performance baseline that combines the network and storage performance over **iSCSI**. This can be a significant help in determining whether your current or future performance is within normal expectations. Similarly, you can use [fio](#) on the *nix side from a target system to establish a performance baseline that combines the network and storage performance over **NFS**.

Troubleshooting and information gathering questions

These questions are presented roughly in order of priority. The answers to these questions should help narrow down possible causes of poor performance.

- All storage traffic for virtualization Target hosts goes over the network. Have you reviewed [Network Performance Expectations and Troubleshooting?](#)
- Were all of the [storage best practices](#) been applied?
 - In particular, check the IO Operation limit and round-robin settings.
- Was a baseline ever created for this host using DiskSpd or fio? (How does our current performance compare?)
- Gather storage specs from VM/storage administrator:
 - Vendor, Model (EMC VMAX 40k, HP 3PAR StoreServ 7000)
 - IO latency SLO (e.g. 5ms 99.999%)
 - IOPS/GB SLO (e.g. 0.68 IOPS/GB for EMC Gold-1)
 - <https://store.emc.com/us/Product-Family/EMC-VMAX-Products/EMC-VMAX-Cloud-Edition/p/EMC-VMAXCloud>
 - Cache type and size (e.g. FAST cache 768GB)
 - Tier, #Pools; if auto – tiering; relocation schedule (e.g. Gold/Silver/Auto/3 pools/etc)
 - Pool detail: (#) drives, RPM, Type (e.g. Pool1: (20) EFD, (30) 15k SAS, Pool 2: (40) 10k SATA)
 - Connection (XXGb Fibre Channel)
 - Dedicated or Shared pool (how many applications/servers)

Conclusion

If you need further help, please contact Delphix Support or Professional Services to assist in getting the best performance possible from your environment.

Related Topics

- [Best Practices for Storage](#)
- [Network Performance Expectations and Troubleshooting](#)
- [Optimal Storage Configuration Parameters for the Delphix Engine](#)

Storage Performance Test Tool notes - Restricted

In addition to displaying the storage results on screen via CLI or retrieving results from a locally downloaded support bundle, Delphix employees have additional options to retrieve the information via ssh or via a support bundle uploaded to <http://upload.delphix.com>.

Download via SSH

Note the destination folder increments based on the job number. storage-test/1, storage-test/2, storage-test/3. etc.

```
delphix$ ls -l /var/delphix/server/storage-test/1/fio_summary*
rw-r----- 1 delphix staff 2310 Apr 29 17:51 /var/delphix/server/storage-test/1/
fio_summary_full.out
rw-r----- 1 delphix staff 1318 Apr 29 17:51 /var/delphix/server/storage-test/1/
fio_summary_grades.out
rw-r----- 1 delphix staff 793 Apr 29 17:51 /var/delphix/server/storage-test/1/
fio_summary.out
```

Note: versions older than 4.2 default to /opt/delphix/monitor/htdocs

Upload to portal; download from portal

1. Upload to portal:
 - a. Login to the Delphix Management application.
 - b. Select System -> Support Logs
 - c. Choose to transfer the logs to Delphix
2. Download steps:
 - a. Go to <http://upload.delphix.com>, login and select "browse uploaded files" to find the file

Extraction

1. Extract the file to locate "storage_tests.tar" in the root of the compressed file.

Architecture for Performance - Hypervisors and Host

The Delphix Engine is a "virtual appliance": a virtual machine guest running within a hypervisor on a physical host. There are a few key best practices we need to keep in mind as we consider this architecture.

Architecture Best Practices Hypervisor Host ESX

Hypervisor

1. ESXi 6.x or 7.0 is recommended. ESXi 5.5 or earlier is no longer supported.
2. HyperThreading (HT) for Intel®-based servers (no HT on AMD CPUs).
 - a. Disable HT in BIOS, on the ESXi Host, **and** disable [HT Sharing](#) on the Delphix VM for consistency. This is our best practice, disable at all levels.
Any other combination may result in a non-deterministic performance.
 - b. When HT cannot be turned off for both the Host and Delphix VM, it should be turned on at all levels, not run in a "mixed-mode".
 - i. HT disablement at a guest level only can result in non-deterministic performance.
 - ii. A dedicated ESXi host, cluster, or DRA is recommended where consistent VDB performance is paramount.
 - c. VM migration to a new host (e.g. VMware HA or vMotion) can create mismatched HT settings.
3. ESXi overhead (resources required for hypervisor cannot be reserved, they must be left unallocated).
 - a. Memory Overhead - 10% of available RAM must not be allocated to guest VMs.
 - *Example: 256GB RAM, allocate 230GB to Delphix VM, leave 26GB for ESX*
 - b. CPU Overhead - At least 2 cores (ideally 4) must not be allocated to guest VMs.
 - *Example: If 16 physical cores are available, allocate 12 to the virtual machines, leaving 4 for the hypervisor*
 - **Why 4?** Certain hypervisor functions require precedence over any virtualized system. If a hypervisor needs more CPU than the amount currently available, it can de-schedule all other virtual processes to ensure adequate CPU resources for the hypervisor. Ensuring the hypervisor will not have to de-schedule any running virtual processes (worlds) by setting aside and not over-subscribing CPUs for virtual functions will leave them available for hypervisor use.
 - c. Even if the Delphix VM is the only VM on a host, the hypervisor is still active and essential; and still needs resources.
4. BIOS Power Management should be set to High Performance where ESXi controls power management.
 - a. It can be impacted by [VMware KB 1018206](#) - poor VM application performance caused by power management settings.
 - b. Ensure that all BIOS-managed C-States other than C0 are disabled if power management is hardware controlled.
 - c. Ensure that all ACPI sleep states above S0 are disabled in the BIOS.
 - d. Examples of popular server lines from Cisco, HP, Dell are below. *Specific models will vary, use the appropriate spec sheet.*
 - i. **UCS:** disable the Processor Power States, disable Power Technology, set Energy Performance to "Performance"
 - ii. **HP ProLiant:** set HP Power Regulator to HP "Static High Performance" mode
 - iii. **Dell:** set BIOS System Profile to "Performance Optimized" mode
5. VMware HA can be enabled; VMware DRS is generally disabled.
6. Blade/Rack Server Firmware and ESXi Drivers should be updated to the latest versions.
7. For Intel®-based servers with E5-2600 v2 processors.
8. Two typical server configurations:
 1. Blade Farm

2. Rack Server
Hyper-Converged configurations are possible for high performance.

Virtual Machine Guest

1. For VM machine settings, see [Virtual Machine Requirements for VMware Platform](#).
2. VMWare Guest Specifications:
 - a. Minimum: 8 vCPU x 64 GB
Small: 8 vCPU x 128GB
Medium: 16 vCPU x 256 GB
Large: 24 vCPU x 512 GB
 - b. Reserve 100% of RAM and CPU:
 - If the ESX host is dedicated to Delphix, CPU and RAM reservations are advised but not necessary, however, swap space will be required on the hypervisor to compensate for the lack of reserved RAM.
 - c. [Hyperthreading](#) - See the ESX host section at the top. Disable HT Sharing on VM, disable HT on ESX Host.
3. Assign single-core sockets for vCPUs in all cases. If there is a compelling reason to use multi-core vCPUs, reference the following article from VMware which describes matching virtual multi-core sockets to the hardware ESX is running on.
[VMware Article on CoresPerSocket](#)
Example: *ESXi Host has 2 socket x 18 core Intel Xeon, Delphix Engine wants 16 vCPU.*
Configure Delphix VM with 2 Virtual Sockets, 8 Cores Per Socket to utilize hardware architecture.
4. Avoid placing other extremely active VMs on the same ESX host.
5. Monitoring - vSphere Threshold Alerts for CPU, Network, Memory, Capacity.
6. To set the number of vCPUs per virtual machine via the vSphere client, please see "[Virtual CPU Configuration](#)" in the Administration guide:
[ESXi 6.0](#)
 - a. [Delphix VM CPU Utilization](#) - Delphix KB article on what makes Delphix VMs similar to other resource-intensive applications
 - b. [Exchange on VMware Best Practices](#) - VMworld 2013 session
 - c. [ESXTOP Reference, Blog](#)
7. Ensure that the latest available VMware drivers and firmware versions are installed for HBAs, NICs, and any other hardware components configured on the Delphix virtual machine. This is a critical step that can have a massive impact on the performance and robustness of our solution.

Related Topics

- [Performance Analytics Management](#)

Target Host OS and Database Configuration Options

This topic describes configuration options to maximize the performance of a target host in a Delphix Engine deployment. These network-tuning changes should improve performance for any data source

OS-Specific Tuning Recommendations

Solaris

When exclusively using Oracle's Direct NFS Feature (dNFS), it is unnecessary to tune the native NFS client. However, tuning network parameters is still relevant and may improve performance.

Tuning the Kernel NFS Client

On systems using Oracle Solaris Zones, the kernel NFS client can only be tuned from the global zone.

On Solaris, by default the maximum I/O size used for NFS read or write requests is 32K. For I/O requests larger than 32K, the I/O is broken down into smaller requests that are serialized. This may result in poor I/O performance. To increase the maximum I/O size:

1. As superuser, add to the `/etc/system` file:

```
* For Delphix: change the maximum NFS block size to 1M
set nfs:nfs3_bsize=0x100000
```

2. Run this command:

```
# echo "nfs3_bsize/W 100000" | mdb -kw
```

It is critical that the above command be executed exactly as shown, with quotations and space. Errors in the command may cause a system panic and reboot.

Tuning TCP Buffer Sizes

On systems using Oracle Solaris Zones, TCP parameters, including buffer sizes, can only be tuned from the global zone or in exclusive-IP non-global zones. Shared-IP non-global zones always inherit TCP parameters from the global zone.

Solaris 10

It is necessary to install a new Service Management Facility (SMF) service that will tune TCP parameters after every boot. These are samples of the files needed to create the service:

File	Installation location
<code>dlpx-tcptune</code>	<code>/lib/svc/method/dlpx-tcptune</code>
<code>dlpx-tune.xml</code>	<code>/var/svc/manifest/site/dlpx-tune.xml</code>

1. As superuser, download the files and install them in the path listed in the **Installation location** in the table.
2. Run the commands:

```
# chmod 755 /lib/svc/method/dlpx-tcptune
# /usr/sbin/svccfg validate /var/svc/manifest/site/dlpx-tune.xml
# /usr/sbin/svccfg import /var/svc/manifest/site/dlpx-tune.xml
# /usr/sbin/svcadm enable site/tcptune
```

Verify that the SMF service ran after being enabled by running the command:

```
# cat `svccprop -p restarter/logfile tcptune`
```

You should see output similar to this:

```
[ May 14 20:02:02 Executing start method ("/lib/svc/method/dlpx-tcptune start"). ]
Tuning TCP Network Parameters
tcp_max_buf adjusted from 1048576 to 16777216
tcp_cwnd_max adjusted from 1048576 to 4194304
tcp_xmit_hiwat adjusted from 49152 to 4194304
tcp_recv_hiwat adjusted from 128000 to 16777216
[ May 14 20:02:02 Method "start" exited with status 0. ]
```

Solaris 11

As superuser

Run the following commands:

```
# ipadm set-prop -p max_buf=16777216 tcp
# ipadm set-prop -p _cwnd_max=4194304 tcp
# ipadm set-prop -p send_buf=4194304 tcp
# ipadm set-prop -p recv_buf=16777216 tcp
```

Linux/Redhat/CentOs

Tuning the Kernel NFS Client

In Linux, the number of simultaneous NFS requests is limited by the Remote Procedure Call (RPC) subsystem. The maximum number of simultaneous requests defaults to 16. Maximize the number of simultaneous requests by changing the kernel tunable `sunrpc.tcp_slot_table_entries` value to 128.

To ensure that the interface does not drop packets because the driver is configured with one receive queue, use the following commands to view the adaptor policy/increase Rx queue length.

```
<LinuxHost> $ ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:22:BB:CC:DD:22
          inet addr:www.xxx.yyy.zzz  Bcast:www.xxx.yyy.zzz  Mask:255.255.255.0
          inet6 addr: feee::222:bbff:fffc:ddd/64  Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:9000  Metric:1
          RX packets:760729910 errors:0 dropped:700 overruns:0 frame:0
          TX packets:309094054 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
```

```
RX bytes:1023150307866 (952.8 GiB) TX bytes:190673864056 (177.5 GiB)
```

RHEL4 through RHEL5.6

1. As superuser, run the following command to change the instantaneous value of simultaneous RPC commands:

```
# sysctl -w sunrpc.tcp_slot_table_entries=128
```

2. Edit the file `/etc/modprobe.d/modprobe.conf.dist` and change the line:

```
install sunrpc /sbin/modprobe --first-time --ignore-install sunrpc && { /bin/
mount -t rpc_pipefs sunrpc /var/lib/nfs/rpc_pipefs > /dev/null 2>&1 || ;;
```

to

```
install sunrpc /sbin/modprobe --first-time --ignore-install sunrpc && { /bin/
mount -t rpc_pipefs sunrpc /var/lib/nfs/rpc_pipefs > /dev/null 2>&1 ; /sbin/
sysctl -w sunrpc.tcp_slot_table_entries=128; }
```

Improper changes to the `modprobe.conf.dist` file may disrupt the use of NFS on the system. Check with your system administrator or operating system vendor for assistance. Save a copy of the `modprobe.conf.dist` in a directory other than `/etc/modprobe.d` before starting.

RHEL 5.7 through RHEL 6.2

1. As superuser, run the following command to change the instantaneous value of simultaneous RPC commands:

```
# sysctl -w sunrpc.tcp_slot_table_entries=128
```

2. If it doesn't already exist, create the file `/etc/modprobe.d/rpcinfo` with the following contents:

```
options sunrpc tcp_slot_table_entries=128
```

RHEL 6.3 onwards

Beginning with RHEL 6.3, the number of RPC slots is dynamically managed by the system and does not need to be tuned. Although the `sunrpc.tcp_slot_table_entries` tuneable still exists, it has a default value of 2, instead of 16 as in prior releases. The maximum number of simultaneous requests is determined by the new tuneable, `sunrpc.tcp_max_slot_table_entries` which has a default value of 65535.

RHEL 6.3 also supports NFSv4.

Tuning TCP Buffer Sizes

Packages should install their configuration files in `/usr/lib/(distribution packages)` or `/usr/local/lib/(local installs)`. Files in `/etc/` are reserved for the local administrator, who may use this logic to override the configuration files installed by vendor packages. It is recommended to prefix all filenames with a two-digit number and a dash, to simplify the ordering of the files. The following is an example approach and should be tested beforehand.

```
echo "Target Kernel Parameter Tunings. This is optional but highly recommended."
echo
```

```

echo "Tuning TCP Buffer Sizes - Parameters should be as below"
echo
echo "This script takes the recommended vendor approach of creating a file in /usr/
lib/sysctl.d and running \"sysctl -p\""
echo "This script will comment out identical settings in /etc/sysctl.conf, which
would otherwise override the Delphix settings"
echo "Admins can at their discretion set larger values, either in sysctl.conf or /
usr/lib/sysctl.d/60-sysctl.conf"
echo "----"
echo "net.ipv4.tcp_timestamps = 1"
echo "net.ipv4.tcp_sack = 1"
echo "net.ipv4.tcp_window_scaling = 1"
echo "net.ipv4.tcp_rmem = 4096 16777216 16777216"
echo "net.ipv4.tcp_wmem = 4096 4194304 16777216"
cat /dev/null > /usr/lib/sysctl.d/60-sysctl.conf
#Set NOW
NOW=$(date +"%m%d%Y%H%M%S")
echo "net.ipv4.tcp_timestamps = 1" >> /usr/lib/sysctl.d/60-sysctl.conf
echo "net.ipv4.tcp_sack = 1" >> /usr/lib/sysctl.d/60-sysctl.conf
echo "net.ipv4.tcp_window_scaling = 1" >> /usr/lib/sysctl.d/60-sysctl.conf
echo "net.ipv4.tcp_rmem = 4096 16777216 16777216" >> /usr/lib/sysctl.d/60-sysctl.conf
echo "net.ipv4.tcp_wmem = 4096 4194304 16777216" >> /usr/lib/sysctl.d/60-sysctl.conf
echo
echo "Running /sbin/sysctl -p /usr/lib/sysctl.d/60-sysctl.conf"
/sbin/sysctl -p /usr/lib/sysctl.d/60-sysctl.conf
if [ $? -ne 0 ]; then
    echo "Command \"sysctl -p /usr/lib/sysctl.d/60-sysctl.conf\" failed; aborting..."
    exit 1
fi
# Make a backup and Comment out similar lines in /etc/sysctl.conf if they exist
# sed will search for lines which are NOT comments ( ^[^\#]* means starts with
anything other than a "#" ) but contain the parameters either with "." or "/"
notation, and add a Comment.
cp /etc/sysctl.conf /tmp/sysctl.conf.$NOW
sed -i "/^[^\#]*net[\./]ipv4[\./]tcp_timestamps/s/^\#Commented out by Delphix /" /etc/
sysctl.conf
sed -i "/^[^\#]*net[\./]ipv4[\./]tcp_sack/s/^\#Commented out by Delphix /" /etc/
sysctl.conf
sed -i "/^[^\#]*net[\./]ipv4[\./]tcp_window_scaling/s/^\#Commented out by Delphix /" /
etc/sysctl.conf
sed -i "/^[^\#]*net[\./]ipv4[\./]tcp_rmem/s/^\#Commented out by Delphix /" /etc/
sysctl.conf
sed -i "/^[^\#]*net[\./]ipv4[\./]tcp_wmem/s/^\#Commented out by Delphix /" /etc/
sysctl.conf
echo
echo "Running /sbin/sysctl -p /etc/sysctl.conf"
/sbin/sysctl -p /etc/sysctl.conf
if [ $? -ne 0 ]; then
    echo "Command \"sysctl -p /etc/sysctl.conf\" failed; aborting..."
    echo "We put a copy of the original at /tmp/sysctl.conf.$NOW"
    exit 1
fi

```

NFSv4 Only - Enabling Recover Lost Locks

RHEL 6.6 onwards

By default, the Redhat NFSv4 client does not attempt to reclaim locks that were lost due to a lease expiration event. This can cause an application to encounter unexpected EIO errors on system calls such as write. The Delphix use case requires that the NFSv4 client attempt to reclaim lost locks that were due to lease expiration. An NFS client module parameter, ' `recover_lost_locks` ', is used to change the default behavior. Use the following command to check if the "recover_lost_locks" option is set to 1:

```
grep recover_lost_locks /etc/modprobe.d/*.conf
```

If the option is currently set to 0, change it to 1. If it is missing, proceed to the next paragraph where instruction is provided on how to add the option.

As superuser, run the following two commands to enable the NFS client to recover the lost locks feature:

```
# cat > /etc/modprobe.d/nfs4-locks.conf <<EOF
options nfs recover_lost_locks=1
EOF

# [ -d "/sys/module/nfs" ] && echo Y > /sys/module/nfs/parameters/recover_lost_locks
```

IBM AIX®

AIX NFSv4 Configuration Requirements (7.1 and 7.2)

1. An NFS Domain must be configured.
2. The `nfsrgyd` service must be running.
3. The NFS server IP address from the Delphix Engine must be mappable to an FQDN.

Configure the nfsv4 domain on the AIX target host

```
bash-3.2# chnfsdom test.com
```

Start the nfsrgyd service and confirm it is active

```
bash-3.2# startsrc -s nfsrgyd
```

```
bash-3.2# lssrc -s nfsrgyd
```

Subsystem	Group	PID	Status
nfsrgyd	nfs	7536760	active

Confirm that IP address can be resolved

```
bash-3.2$ host 172.16.105.81
81.105.16.172.in-addr.arpa is dc011.delphix.com
```

Reference: [IBM AIX: HOW TO SETUP NFSV4 MOUNT IN CLIENT AND SERVER](#)

Tuning the Kernel NFS Client

On AIX, by default the maximum I/O size used for NFS read or write requests is 64K. When Oracle does I/O larger than 64K, the I/O is broken down into smaller requests that are serialized. This may result in poor I/O performance. IBM can provide an Authorized Program Analysis Report (APAR) that allows the I/O size to be configured to a larger value.

1. Determine the appropriate APAR for the version of AIX you are using:

AIX Version	APAR Name
6.1	IV24594
7.1	IV24688

2. Check if the required APAR is already installed by running this command:

```
# /usr/sbin/instfix -ik IV24594
```

If the APAR is installed, you will see a message similar to this:

```
All filesets for IV24594 were found.
```

If the APAR is not yet installed, you will see a message similar to this:

```
There was no data for IV24594 in the fix database.
```

3. Download and install the APAR, as necessary. To find the APARs, use the main search function at <http://www.ibm.com/us/en/>, specifying the name of the APAR you are looking for from step 1. A system reboot is necessary after installing the APAR
4. Configure the maximum read and write sizes using the commands below:

```
# nfsd -p -o nfs_max_read_size=524288
# nfsd -p -o nfs_max_write_size=524288
```

5. Confirm the correct settings using the command:

```
# nfsd -L nfs_max_read_size -L nfs_max_write_size
```

You should see an output similar to this:

```

NAME CUR DEF BOOT MIN MAX UNIT TYPE
DEPENDENCIES
-----
nfs_max_read_size 512K 64K 512K 512 512K Bytes D
-----
nfs_max_write_size 512K 64K 512K 512 512K Bytes D
-----

```

Tuning Delayed TCP Acknowledgements

By default, AIX implements a 200ms delay for TCP acknowledgments. However, it has been found that disabling this behavior can provide significant performance improvements.

To disable delayed ACKs on AIX the following command can be used:

```
# /usr/sbin/no -o tcp_nodelayack=1
```

To make the change permanent use:

```
# /usr/sbin/no -p -o tcp_nodelayack=1
```

HP-UX

Tuning the Kernel NFS Client

On HP-UX, by default the maximum I/O size used for NFS read or write requests is 32K. For I/O requests larger than 32K, the I/O is broken down into smaller requests that are serialized. This may result in poor I/O performance.

1. As superuser, run the following command:

```
# /usr/sbin/kctune nfs3_bsize=1048576
```

2. Confirm the changes have occurred and are persistent by running the following command and checking the output:

```
# grep nfs3 /stand/system
tunable nfs3_bsize      1048576
```

Tuning TCP Buffer Sizes

1. As superuser, edit the `/etc/rc.config.d/nddconf` file, adding or replacing the following entries:

```

TRANSPORT_NAME[0]=tcp
NDD_NAME[0]=tcp_recv_hiwater_def
NDD_VALUE[0]=16777216
#
TRANSPORT_NAME[1]=tcp
NDD_NAME[1]=tcp_xmit_hiwater_def
NDD_VALUE[1]=4194304

```

In this example, the array indices are shown as `0` and `1`. In the actual configuration file, each index used must be strictly increasing, with no missing entries. See the comments at the beginning of `/etc/rc.config.d/nddconf` for more information.

2. Run the command:

```
/usr/bin/ndd -c
```

3. Confirm the settings:

```
# ndd -get /dev/tcp tcp_recv_hiwater_def
16777216
# ndd -get /dev/tcp tcp_xmit_hiwater_def
4194304
```

OS-Specific Tuning Recommendations for Windows

iSCSI Tuning

These are our recommendations for the Windows iSCSI initiator configuration. Please note that the parameters below will affect **all** applications running on the Windows target host, so you should make sure that the following recommendations do not contradict best practices for other applications running on the host.

For details about the Windows iSCSI configuration, see [Requirements for Windows iSCSI Configuration](#).

For targets running Windows Server, the iSCSI initiator driver timers can be found at: `HKLM\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\<Parameters>`. Please see [How to Modify the Windows Registry](#) on the Microsoft Support site for details about configuring registry settings.

Registry Value	Type	Default	Recommended	Comments
MaxTransferLength	REG_DWORD	262144	131072	This controls the maximum data size of an I/O request. A value of 128K is optimal for the Delphix Engine as it reduces the segmentation of the packets as they go through the stack.
MaxBurstLength	REG_DWORD	262144	131072	This is the negotiated maximum burst length. 128K is the optimal size for the Delphix Engine.
MaxPendingRequests	REG_DWORD	255	512	This setting controls the maximum number of outstanding requests allowed by the initiator. At most this many requests will be sent to the target before receiving a response for any of the requests.

Registry Value	Type	Default	Recommended	Comments
MaxRecvDataSegmentLength	REG_DWORD	65536	131072	This is the negotiated MaxRecvDataSegmentLength.

Receive Side Scaling

Follow the instructions here to enable RSS: [Enable Receive Side Scaling \(RSS\) on Staging/Target Network Interfaces](#) hosted on the target, such as Oracle, SQL Server, Sybase, or vFiles. They should be applied to all Delphix targets.

Related Topics

- [Best Practices for Target DB and OS Settings](#)
- [Performance Analytics Management](#)

Usage Data Management

The Delphix User-click Analytics feature is a lightweight method to capture how users interact with Delphix product user interfaces. The goal of capturing this data is to get a better understanding of product usage, engagement, and user behavior, and to use this data to improve Delphix products and customer experience. This feature is enabled by default for customers deploying on or upgrading to this version. User-click Analytics may also be disabled via the UI.

Disabling User-click Analytics

⚠ This procedure will disable user-click analytics on both the Delphix Engine and Delphix Self-Service.


Procedure to disable user-click analytics via the GUI

Complete the following steps to disable user-click analytics via the GUI.

1. Login to the Delphix Setup application using the sysadmin **username** and **password**.

2. From the Outbound Connectivity panel, click **Modify**.

3. In the **Outbound Connectivity** window, uncheck **Enable Usage Analytics**.



Outbound Connectivity

WEB PROXY
The Web Proxy Server will be used to communicate with Delphix Corp. for support, troubleshooting, upgrades, updates, and patches.

Configure web proxy

PHONE HOME SERVICE
If enabled, this service will automatically send a minimal support bundle once a day to the Delphix support site over HTTPS. This will help with future support and troubleshooting. A connection to the internet, either directly or via web proxy is required.

Enable phone home service

USER-CLICK ANALYTICS
If enabled, this service will automatically send a stream of anonymous, non-personal metadata describing user interaction with the product's user interface. This data will help us to better understand how our products are being used, and to improve our products and services.

Enable Usage Analytics

SMTP
Configure the Delphix Engine's SMTP sending service to enable email notifications. Your sysadmin email will be used for receiving system reports, events, and fault notifications.

Use an existing SMTP server
Enable email notifications for faults and events

⚠ SMTP not configured - functionality reduced: Not able to send or receive events and notifications.

Cancel Save

4. Click **Save**.

Related Topics

- [Configuration](#)

Starting, Stopping, and Restarting Your Engine

Several options are available to start, stop, restart, or reset the Delphix engine for maintenance or debugging purposes.

Factory Reset


This option returns the Delphix Engine to the "factory default" state. This removes all the data and configurations except for the network configuration and NTP settings. The network configuration and NTP settings are retained in order to make it is easy for a user to start the Delphix Engine otherwise you will need to perform the initial steps to configure the network via the VMware Hypervisor console.

It is recommended to shut down and remove all VDBs before resetting the Delphix Engine. Failure to do so could lead to stale data mounts in target environments. For the same reason, disable all dSources that use validated sync.

Use Factory Reset only when a complete reset and reconfiguration of the Delphix Engine is necessary, as all Delphix Engine objects will be de-allocated.

Complete the following steps to reset the Delphix Engine via GUI:

1. Connect to the Delphix Setup application and log in as a system administrator.
2. From the **Dashboard** panel, click the Actions menu (...) on the top right and select **Factory Reset**.



3. Click **OK**.

Restart

If the Delphix Java process is in a bad state, you may need to restart it to get it operational again.

You can restart the management process safely without shutting down VDBs or disabling dSources because restarting this process has no impact on running VDBs or dSources.

Complete the following steps to restart the Delphix Management Process via GUI:

1. Login to the Delphix Setup application as a system administrator.

- From the **Dashboard** panel, click the Actions menu (...) on the top right and select **Restart**.



- Click **OK**.

Complete the following steps to restart the Delphix Management Process via CLI:

- Log in to the CLI as a system administrator.
- Go to `system > restart`
- `delphix system restart *> commit`


Reboot

If the management stack for your Delphix Engine hangs, you will need to perform a reboot.

Before performing a reboot, all your VDBs must be shut down and dSources disabled to maintain data integrity.

Complete the following steps to reboot the Delphix Engine via GUI:

- Login to the Delphix Setup application as a system administrator.
- From the **Dashboard** panel, click the Actions menu (...) on the top right and select **Reboot**.



- Click **OK**.

Complete the following steps to reboot the Delphix Engine via CLI:

1. Log in to the CLI as a system administrator.
2. Go to system > reboot
3. delphix system reboot *> commit

i Reboot vs Restart

The only difference between the Restart and Reboot is that the Restart functions without shutting down the VDBs and disabling the dSources.


Power Off

Occasionally, it is necessary to shut down the Delphix Engine for maintenance purposes. Before performing a shutdown, all your VDBs must be shut down and dSources disabled to maintain data integrity.

Once the Delphix Engine is powered off, you can power it back on through your hypervisor console.

Use the following steps to shut down the Delphix Engine via GUI:

1. Log in to the Delphix Setup application as a system administrator.
2. From the **Dashboard** panel, click the Actions menu (...) on the top right and select **Power Off**.



3. Click **OK**.

Use the following steps to shut down the Delphix Engine via CLI:

1. ssh sysadmin@yourengine
2. delphix > system
3. delphix system > shutdown
4. delphix system shutdown *> commit

Related Topics

- [Initial Setup](#)
- [Customizing the Delphix Engine System Settings](#)
- [Configuration](#)

Introduction to Privilege Elevation Profiles

This topic introduces the concept of Privilege Elevation Profiles, how they are managed, and how they are supported. Privilege Elevation Profiles exist to provide the Delphix Engine with a mechanism for running privileged commands in a secure way to achieve the following:

- Mount and Unmount NFS filesystems
- Create and Remove directories in paths not owned by the Delphix OS user
- Examine the running process list
- Run commands as root

Privilege Elevation Profiles is an advanced CLI topic and are not documented as part of the general Delphix Engine User Guide. Changes to the default sudo-based profile scripts, or the creation of new profiles that do not work as expected, can cause serious problems and render the Delphix Engine unusable. This article is aimed at advanced end-users and Delphix Professional Services consultants.

Support for Privilege Elevation Profiles

 Writing and troubleshooting scripts, such as those required for Privilege Elevation Profiles, is out of scope and not covered by Delphix Support.

Privilege Elevation Profiles need to be tailor-made to work with non-standard environments that may use third-party or proprietary privilege elevation mechanisms other than sudo. Customers are strongly encouraged to work with Delphix Professional Services to formulate reliable profile scripts. There is nothing that prevents customers from creating their own profile scripts. However, customers bear full responsibility for supporting and troubleshooting their own profile scripts. Support for profile scripts created by our Professional Services consultants is still supported by Professional Services.

How do Privilege Elevation Profiles Work?

Privilege Elevation Profiles exist within a two-tier cascading hierarchy. This means there is one default profile for the entire Delphix Engine that should contain scripts for all the operations that require privilege elevation. Additional profiles may contain a subset of the scripts. When a non-default profile is used, the Delphix Engine uses that profile's scripts where they exist and reverts to the scripts in the default profile if no script for the operation exists. By default, the Delphix Engine ships with simple scripts that pass commands to the standard UNIX **sudo** command.

All Environments added to the Delphix Engine get added with the default Privilege Elevation Profile. The profile can be assigned on a per-host basis. Below shows how a host using a non-standard profile will use scripts in the cascading model.

default profile (sudo)	custom profile (myProfile)	host profile	script used
dlpx_mount	my_mount	myProfile	my_mount
dlpx_umount	my_umount		my_umount
dlpx_rmdir			dlpx_rmdir

default profile (sudo)	custom profile (myProfile)	host profile	script used
dlpx_mkdir			dlpx_mkdir
dlpx_ps			dlpx_ps
dlpx_pfexec			dlpx_pfexec

Related Topics

- [Privilege Elevation Profiles and Delphix Replication](#)
- [Installation and Initial System Configuration](#)

Upgrade

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

 Unknown macro: 'navbox'

Upgrading the Delphix Engine: Overview

Introduction

Upgrading the Delphix Engine is a multi-step process that requires some preparation. The engine upgrade process will affect the availability of the Delphix Engine administrative interface and virtual datasets during the operation based on the type of upgrade chosen. For large configurations and Full/Apply Now upgrades, it can take one to two hours; Delay the Reboot upgrades are typically complete within 15 to 30 minutes. Please refer to [Upgrade Matrix](#) and the **Version Compatibility and Support Pre-Checks** callout below before proceeding with an upgrade.

The following sections explain the steps involved in the upgrade process with links to detailed instructions for proceeding through each of them.

Types of Upgrade

There are two types of upgrades, which are characterized by the impact they have on VDBs during the operation:

 There is no need to reboot unless instructed to receive a specific fix from Delphix.

Upgrade Type	Description
Delay the Reboot	The user interface, API, and CLI (Command Line Interface) will only be available to the user performing the upgrade. dSources will stop refreshing from production. Policies execution will be delayed until after the upgrade has been completed. Jobs will be canceled (and resumed after upgrade if supported). Access to VDB data will not be affected by this upgrade and can be used normally.
Apply Now	In addition to performing an application upgrade, DelphixOS, the operating system that runs Delphix, will be upgraded and the machine will reboot to the new OS as part of the upgrade process. The Delphix Engine will automatically disable all VDBs and dSources during the upgrade process in order to safely reboot to the new version, and thus you should schedule downtime for your VDB applications.

Outline of the Upgrade Process

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

The following is an outline of the steps for upgrading the Delphix Engine:

1. [Upload](#) the upgrade image to the Delphix Engine.
2. [Verify](#) and [resolve](#) the system requirements and known defects before starting the upgrade.

3. Schedule the appropriate downtime.
4. [Start](#) the upgrade and choose the upgrade type.
5. [Address](#) any runtime failures that happen as part of the upgrade.
6. [Verify](#) that the upgrade was completed successfully.

Related Topics

- [Upgrade](#)

Upgrade Prerequisites

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

Scheduling Downtime

To determine if an upgrade will require a reboot and VDB downtime refer to [Upgrade Matrix](#).

If the OS will not be updated as part of the upgrade, then the upgrade process will have no impact on the availability of VDBs, and you do not need to schedule any downtime for your VDB applications.

If the OS will be updated as part of the upgrade, then you should schedule appropriate downtime for your VDB applications. The Delphix Engine will automatically disable VDBs and dSources during the upgrade. The length of the downtime will be proportional to the number of VDBs.

Long-running jobs including replication and SnapSync will fail during any upgrade.

The upgrade file for the version to which you want to upgrade should be downloaded from the [Delphix download site](#). From 6.0.2.0 onwards, the upgrade images are packaged with the latest version of upgrade pre-checks. Please use the latest upgrade image from the download site since the checks will be updated on a need-by basis.

Delphix Upgrade images are approximately 5GB in size; it is recommended to have both a fast internet connection to the Delphix download site as well as to the Delphix Engine.

The upgrade image should be downloaded or moved to a location accessible to the computer used for navigating the Delphix Management application.

Upgrading Replication Source

Delphix Engines can only perform replication to engines on the same or higher version. Upgrading a replication source without upgrading its replication targets could cause replication between those peers to fail.

Verifying the Integrity of the Downloaded Upgrade Image

SHA256 checksums can be used to verify the integrity of files downloaded via download.delphix.com.

The sha256sum checksum listed next to each download will help ensure that a file has not changed as a result of a faulty file transfer, a disk error, or non-malicious meddling.

To calculate the checksum of downloaded files and compare them to the checksum listed on the download portal, the following utilities can be used:

- On Windows, `certutil -HashFile <upgrade.tar.gz> SHA256`
- On macOS or Unix servers with Perl installed, `shasum -a 256 <upgrade.tar.gz>`
- On Unix servers with `coreutils` installed, `sha256sum <upgrade.tar.gz>`.

Verifying Connectivity to Datasets and Environments

For Upgrades from versions > 5.3.6.0 and <6.0.0.0, your Delphix Engine will automatically quiesce all VDBs and dSources during the upgrade process in order to safely reboot to the new version.

For Upgrades from versions >= 6.0.0.0, based on the [Upgrade Matrix](#) you can choose to postpone the VDB downtime and update just the running stack and related packages by choosing "Delay the Restart" option. You can update the OS at a later date by completing the upgrade.

At the end of the upgrade process, the Delphix Engine will also update the Delphix platform toolkit on each connected environment.


To perform these tasks, the Delphix Engine must be able to connect to the environments in which datasets exist and must have credentials to connect to datasets or applications. Environments involved must be properly configured to enable script execution.

When an environment, dataset, or application is unreachable or misconfigured, the upgrade may encounter access errors, and may not be able to re-start it after the upgrade has been applied.

If you know that a dataset is unavailable for some reason (e.g. host not reachable on the network, or the hosting server is down), it is recommended that you **disable** it (if necessary utilize **Force Disable**).

The **Impact of Upgrade** section in the **Details** tab contains the list of datasets that have been identified as unreachable, misconfigured or not behaving correctly and are at risk of not functioning after the upgrade.

Review the warnings in this section (if any) and take appropriate actions before applying the upgrade.



Related Topics

- [Upgrading the Delphix Engine: Overview](#)
- [Upgrade Matrix](#)

Downloading the Upgrade Image

The upgrade file for the version to which you want to upgrade should be downloaded from the [Delphix download site](#). From 6.0.2.0 onwards, the upgrade images are packaged with the latest version of upgrade pre-checks. Please use the latest upgrade image from the download site since the checks will be updated on a need-by basis.

Delphix Upgrade images are approximately 5GB in size; it is recommended to have both a fast internet connection to the Delphix download site as well as to the Delphix Engine.

The upgrade image should be downloaded or moved to a location accessible to the computer used for navigating the Delphix Management application.

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

For upgrades from versions > 5.3.6.0 and < 6.0.0.0, you will download a migration image that is specific to the virtualization platform that your Delphix Engine is running on. The platforms for which images are provided include:

- Amazon AWS
- Microsoft Azure
- VMware ESX

For example, an ESX migration image will have a filename

like `Delphix_6.0.0.0_2019-12-06-09-07_Standard_ESX_Migration.tar.gz`.

For upgrades from versions $\geq 6.0.0.0$, you will download an upgrade image that is generic to all the virtualization platforms that your Delphix Engine is running on. The Upgrade image will have a filename like

`Delphix_6.0.1.0_2020-03-17-00-39_Standard_Upgrade.tar`.

SHA256 checksums, which can be used to verify the integrity of any downloaded files, are available via download.delphix.com. Each download link should list a `sha256sum` checksum underneath. Additionally, newer releases may also contain a downloadable SHA256SUMS file, containing the checksums of all files in that directory.

To calculate the checksum of downloaded files, and confirm that they match those in the SHA256SUMS file, the following utilities can be used:

- On Windows, `certutil -HashFile <upgrade.tar.gz> SHA256`
- On macOS or Unix servers with Perl installed, `shasum -a 256 <upgrade.tar.gz>`
- On Unix servers with `coreutils` installed, `sha256sum <upgrade.tar.gz>`.

Related Topics

- [Upgrade Prerequisites](#)
- [Upgrade Verification](#)

Uploading the Upgrade Image

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS


Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

The procedure for uploading an upgrade version to the Delphix Engine is:

1. Login to the **Delphix Setup** application.
2. In the **Software Version** panel, click **View**.




The screenshot shows the 'Dashboard' interface. At the top, there is a 'Software Version' panel with a 'View' link. Below this, the 'Engine' section displays the following information:

- Current Version:** Dynamic Data Platform 6.1.0.0-snapshot.20210913234913859+jenkins-ops-devops-gate-master-appliance-build-master-post-push-6092
- Build Date:** Sep 14, 2021 5:13:13 AM
- Latest Version:** 6.1.0.0-snapshot.20210913234913859+jenkins-ops-devops-gate-master-appliance-build-master-post-push-6092


Below the Engine section, there is a 'Data Control Tower Connector' section with an 'Enable' button.

3. Click the **plus** icon to upload a new version.
4. Select the upgrade version you downloaded from the download site.



The upload begins on the drop or selection of a file (upgrade image) into the **Upload Upgrade Image** window. You will be able to view the progress in the **Upload Status** area.

5. In the **Upload Complete** dialog, click **OK**. You will see the new version appear on the page.



Once the upgrade image upload is complete, upgrade verification will start automatically

This verification process will notify you of any potential problems that require intervention from either you or Delphix support. The new version of the Delphix Engine cannot be installed before verification has completed. You can track the progress of verification from the **Actions** sidebar. If issues are found during the verification process, they will be listed on the version page. An upgrade will not be possible until the issues are resolved.

For more details, see [Resolving Upgrade Checks](#).

Related Topics

- [Resolving Upgrade Checks](#)

- [Upgrading the Delphix Engine: Overview](#)

Upgrade Verification

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

The Delphix Engine provides a feature that allows you to verify an upgrade before applying it. Each new version of the Delphix Engine can introduce new requirements for networking, hypervisor usage, the configuration of Delphix Engine objects and more. Verification can determine if the current Delphix Engine does not meet these requirements. It can also detect if the Delphix Engine is in an unexpected state before an upgrade, in which case the best solution might be to contact support. It is recommended that customers read the [release notes](#) for any version they are upgrading to and look for new Delphix Engine requirements.

As soon as the upgrade image is [uploaded](#), a verification job is started. This verification job will use information stored within the uploaded upgrade image to examine the state of the Delphix Engine and make a best effort to validate that applying the upgrade image will succeed. It also will notify the user of any potential problems that require either customer or support intervention.

After resolving problems noted by the verification process, verification can be run again. It is expected that customers will continue to fix problems and re-verify until no more problems remain, or until none of the remaining problems are critical.


The verification does a “dry run” of some of the upgrade procedures in order to alert the administrator of potential problems before continuing with the upgrade. It is strongly recommended that you perform this verification a day or two in advance, before your upgrade downtime begins, in order to give yourself time to address any problems flagged by the verification. Perform a re-verify closer to the upgrade, when issues have been resolved.

The procedure for verifying an upgrade is:


1. Login to the **Delphix Setup** application.
2. In the **System Upgrade Management** panel, click **View**.
3. On the left-hand side, select the version to which you will be upgrading. Details on the version will be displayed on the right.
4. In the upper right-hand section of the Details tab, click **Verify Upgrade**. Verification will run in the background. You can view the different stages with the progress during the upgrade process in the **Action** sidebar.
5. Click the **Verification Results** option to view problems during verification.
6. Click the **Verification Steps** option to view verification steps in detail.

Understanding the Verification Page

This section provides details about various components on the **Verification** page.



- 1 **Verification Package Version** displays the version of verification checks that are bundled with an upgrade version.
- 2 **Verification Results** provides result details about the verification outcomes. This option is selected by default. The verification result detail includes upgrade severity, duration, and reason
- 3 **Hide Resolved/Ignored** enables hiding of all verification items that are Resolved or Ignored.
- 4 **Ignore All Warnings** allows you to ignore all warning items during the upgrade process.



5 Verification Steps allow you to view a complete list of verification actions, status, and durations.

Verification
Report

Verification Package Version
0.0.0

Verification Results
 Verification Steps

Description	Status	Duration
Verify Storage available	PASS	00:00:00
Verify integrity of internal file systems	PASS	00:00:00
Initialize verification	PASS	00:04:15
Verify no differences in SSL/TLS configuration between Masking and Virtualization engines	PASS	00:00:00
Performing internal database migration	PASS	00:00:01
Verify post-upgrade compatibility of Delphix Replication target	PASS	00:00:00
Verify Delphix filesystem (DxFS) permissions	PASS	00:00:00
Verify no unsupported ciphers in Virtualization engine	PASS	00:00:00
Verify OS configuration	SKIPPED	00:00:00
Verify correct migration image for platform.	SKIPPED	00:00:00
Verify that the OS configuration can be migrated from illumos to Linux engines	SKIPPED	00:00:00
Verify not using native Postgres.	SKIPPED	00:00:00
Verify not using Cross-platform provisioning (XPP)	SKIPPED	00:00:00
Verify Hotfixes are resolved	▲ FAIL	00:00:00
Verify post-upgrade compatibility of Delphix Replication source	▲ FAIL	00:00:00

1 to 15 of 15
First
Previous
Page 1 of 1
Next
Last

▲ FAIL

DESCRIPTION
Verify Hotfixes are resolved

START TIME
Aug 8, 2019 3:26:41 PM

END TIME
Aug 8, 2019 3:26:41 PM

VERIFICATION RESULTS
Hotfix metadata malformed
Hotfixes not integrated


Related Topics

- [Release Notes](#)
- [Upgrading the Delphix Engine: Overview](#)

Resolving Upgrade Checks

Overview

When verifying an upgrade, the verification job can sometimes detect particular action items called upgrade checks. These checks are items that the Delphix Engine is not capable of fixing on its own, and require customer and/or Delphix support action. Every upgrade check that appears must be resolved or ignored before the upgrade can proceed, with the exception of INFORMATIONAL checks which require no action by users. Below is a workflow for various severity types and upgrade check action.



Upgrade Severities Checks

Severity	Description
CRITICAL	These checks indicate a potential problem post-upgrade with the Delphix Engine. The Delphix Engine cannot be upgraded while CRITICAL checks are present; it is likely that the engine will fail to upgrade or break catastrophically post-upgrade.
WARNING	These checks indicate problems that are local to some objects on the Delphix Engine, but the overall engine can continue running. For example, WARNING checks may indicate a problem with the configuration of a particular Windows environment and indicate that if an upgrade occurs, that environment will not function properly.

Severity	Description
INFORMATIONAL	These checks provide important information about your upgrade but do not require immediate action.

The following is a list of checks that have been added to the Delphix Engine:

1. **OS tunable settings** (CRITICAL) - Upgrade verification checks that only a certain set of operating system tunables have been adjusted. If tunable settings not on this acceptable list have been changed, please contact Delphix support.
2. **Hotfix** (CRITICAL) - Upgrade verification checks if any hotfixes have been applied to the Delphix Engine, and if any exist that are not resolved in the Upgrade version, this check will generate a Critical result. In this scenario, the Engine may still be upgraded, but we ask that you contact Delphix support to assist in addressing the check.
3. **Snapshot directory visible** (CRITICAL) - This is an internal Delphix Engine problem that Delphix support can resolve for our customers.
4. **Replication** (WARNING) - It indicates that your Delphix Engine has replication set up as either a target or source. Delphix Engines can only perform replication to engines on the same version or higher, so it is recommended to upgrade all Engines in a Replication configuration at the same time, depending on the requirements. For more information on forward-compatible replication, refer to [Replication Overview](#).
5. **Storage available** (WARNING) - Upgrade verification checks that there is sufficient space to perform data operations.
6. **SSL/TLS configuration** (WARNING) - Upgrade verification checks that the security configuration of the engine is up to date and properly configured.

An Upgrade check can result in multiple check results. Check results are essentially a to-do list of action items required before performing an upgrade. For example, the OS Tunables will create one line item in the UI for each present tunable that is not on our acceptable list. There are two actions that can be performed on each check result.

Upgrade Check Actions

Action	Description
resolve	You have taken the action necessary to solve the problem and complete the action item. Marking a result as resolved means that the customer believes the check result indicated will no longer be a problem. The next verification run will check the result again, but if it is fixed will not create any new check results.
ignore	You have no intention of fixing the problem indicated by the check result. Future verification runs will not generate that check result again if it is present. CRITICAL severity checks cannot be ignored.


Examples of Upgrade Check Actions

Case A: Verification Results from Critical Severity Check

Consider a case where Delphix Engine previously had a few hotfixes installed by Delphix Support to resolve an issue.

During the upgrade process, a critical severity message is displayed indicating the hotfixes that was used to resolve issue might be lost or reappear after the upgrade.

Hence new upgrade version is missing the root cause fix for that issue. In this example, the action is either to contact Delphix Support or upgrade to a new version that no longer needs these hotfixes. After taking the necessary action, you must mark resolved for the verification steps to complete for a successful upgrade.



1 It is mandatory to resolve all critical severity errors


Case B: Verification Results Warning Severity Check with Ignore All Warnings

Consider a case where Delphix Engine is configured as a Delphix Replication source for another Delphix Engine.

During the upgrade process, a warning message appears because the target engine is on a version that is incompatible with the source engine version.

Hence replication is interrupted until the target engine is upgraded to a newer version.

In this example, the action steps are either to mark the issue as **Mark Resolved** or **Mark Ignored** because a replication-compatibility is acceptable with a few limitations (as indicated in below warning message).




- 1 All Warnings must either be **Mark Resolved** or **Mark Ignored**
- 2 You also can use **Ignore All Warnings** option to disregard all warning messages at once.

Case C: Verification Steps Upgrade Failure

When you want to see more details about the verification process use the **Verification Steps** radio button. Verification Steps is an alternate way to view the verification status. It includes all the steps that were performed by the verification process, the status, description and the duration of the step.

In this example below verification status, **FAIL** indicates that a verification result was generated. In the **Verification Results** view, a **FAIL** status on this screen may correspond to a CRITICAL, WARNING, or INFORMATIONAL severity.



- 1 Failure Severity
- 2 The description indicates why upgrade failure occurred
- 3 Verification Results

Related Topics

- [Upgrade Verification](#)
- [Uploading the Upgrade Image](#)

Applying the Upgrade

VERSION COMPATIBILITY AND SUPPORT PRE-CHECKS

Please refer to [Upgrade Matrix](#) before upgrading to a newer version.

Customers running version 5.3.9 and earlier that are requesting an upgrade to 6.0.0.0 and above, please contact Delphix Support to help coordinate this upgrade.

Upgrading from 6.0.x to 6.0.x includes pre-checks packaged in the upgrade image, contacting Delphix Support for this upgrade is **not required** (e.g. 6.0.0.0 → 6.0.9.0).

Once you have uploaded an upgrade version, verified the upgrade, optionally reviewed the warnings in the Impact of Upgrade section, scheduled downtime pertaining to the type of upgrade you are performing, you can apply the upgrade.

1. Login to the **Delphix Setup** application.
2. In the **Upgrade Images** panel, click **View**.
3. On the left-hand side, select the version to which you will be upgrading.
4. Click **Apply Upgrade** to initiate the upgrade process.
 - For upgrades on engines already running 6.0.0.0 and greater, the user can choose the type of upgrade they want to perform.

The upgrade will run in the background. You can view the progress of the upgrade in the Action sidebar. Only the current system admin user can view the progress.

The status of the upgrade will be visible on the screen - if the upgrade is successful, the page will be redirected to the login view.

If an **Apply Now** upgrade fails, the appliance will automatically roll back to the version running prior to the upgrade.

The version page will show the new version in an UPLOADED state and the Action sidebar will show that a rollback was performed. If automatic rollback was disabled through the CLI (not advised), you will have to contact support to proceed further, since you may not even be able to log in to the Delphix Engine.

Failure to Quiesce a Dataset

If Upgrade has failed to quiesce a dataset, it will pause and you will see the banner at the top of the **Upgrade** page as shown below:

While the upgrade is paused, datasets that have been quiesced are unavailable until you either roll back or continue the upgrade.

To review the list of failures, open the **Report** tab:

Upgrade Image - 6.0.10.1

Cancel Upgrade
Retry Upgrade
Continue Upgrade

Upgrade is unable to quiesce some datasets and has paused.

Version 6.0.10.1 - Paused on failures

Release Date Oct 29, 2021	Verify Date Not verified
OS Version 6.0.10.1	Minimum OS Version 6.0.4.0

Verification
Report

The datasets listed below were identified as having issues preventing their quiescence, and may not be available after the upgrade is complete.

You can attempt to remedy the issues by:

- Clicking Continue Upgrade (those datasets may not restart after the upgrade)
- Going to the Admin Application in a different browser (e.g. Chrome or IE) or using an incognito window in this browser and resolving issues such as wrong passwords, or stopping the datasets (using Force Disable), then clicking Retry

If you choose to Cancel Upgrade, all changes will be reversed, upgrade will end, and the Engine will be in the state it was in before you started the upgrade.

Filter: none

Info	Source	Hosts	Failure	Action
	VDBOMSR94FDE4_JYH	ac12201.dcol1.delphix.com	Stress option upg.quiesce...	Disable stress options.
	VDBOMSR94FDE4_JYH	ac12201.dcol1.delphix.com	UPGRADE_QUIESCE_SOUR...	Try the job again.

The datasets listed in the report were identified as having issues that prevented them from being quiesced, and may not be available after the upgrade is complete. Review the messages in the report and take the suggested corrective actions.

If you think that the errors may be the result of transient failures, you can click **Retry Upgrade** to try again. Otherwise, it is recommended that you manually quiesce datasets that are still running. To do so:

1. Use a different browser or use an incognito window to go to the **Delphix Management** application.

Failure to do so might result in losing the current session and you may not be able to return to the Upgrade pause screen.

2. Either resolve issues such as a wrong password or stop the dataset using **Force Disable**.
3. In the original browser or window, click **Retry Upgrade** to try applying the upgrade again.

If you want to ignore the failures to quiesce datasets and proceed with the upgrade:

1. Click **Continue Upgrade**. This will attempt to quiesce all datasets which have not yet been quiesced, but will not pause on failures.

 This may result in datasets remaining unavailable after the upgrade is complete and the Delphix Engine restarts, since the underlying storage that backs the datasets will be unreachable during the upgrade. This may cause the databases or applications to failover or transition to a failure state, thus requiring administrator intervention to recover.

2. Review the messages in the report and take the suggested corrective actions.
3. If any of the listed datasets are critical, and you are unable to resolve the configuration errors in the report, you can **Rollback** the upgrade. If you choose **Rollback**, all changes will be reversed, the upgrade will end, and the Delphix Engine will be in the state it was in before you started the upgrade.

Related Topics

- [Upgrading the Delphix Engine: Overview](#)

Post Upgrade

After the upgrade is done, you will be redirected back to the login page.

Login to **Delphix Setup** to make sure that the upgrade succeeded and that the new version is running.

A **post-upgrade cleanup** job is run automatically after the upgrade is done. This job refreshes environments and re-enables sources to bring objects back into a working state. VDB access will not be available at all until the environments have been refreshed and the objects are re-enabled by the job **Perform cleanup tasks following a Delphix Engine upgrade**.

Failures

If you used **Continue Upgrade** to force the upgrade as is described in [Failure to Quiesce a Dataset](#), the **Report** tab will contain the list of pre-upgrade quiesce failures. Upgrade will make a best effort to restore functionality to these datasets, but it may still hit the same errors that prevented the datasets from being quiesced successfully before the upgrade. You may need to manually bring up these datasets on the target hosts.

Related Topics

- [Upgrade](#)
- [Upgrading the Delphix Engine: Overview](#)
- [Uploading the Upgrade Image](#)
- [Applying the Upgrade](#)

Security

 Unknown macro: 'navbox'

Security Principles

The Delphix approach is based on:

Embrace separation of duties: Isolate and compartmentalize capabilities and privileges and never give or concentrate access to a single role.

Apply the principle of least privilege: Users should obtain only those privileges needed to do their jobs and only for as long as they are needed.

Use an open, simple design: Make security mechanisms simple and easy to use, and rely on proven, peer-reviewed solutions.

Use a layered defense: Provide no single point of failure; if one layer fails to catch an error, catch it in another layer.

Use complete mediation and authentication: Control and check every access point every time.

Use fail-safes: Deny access when not explicitly authorized. Prevent faults from causing an opportunity to compromise.

Protect data at rest and data in motion: Utilize common security protocols as well as features of the source database and database software to protect data at all times.

Minimize the attack surface: Present the minimum sockets, services, webpages, and accounts necessary to operate.

Don't rely on obscurity: Be secure even if everything but the key is known.

Audit and monitor everything: Provide a tamper-proof trail of evidence.

Leverage the environment: Design the Delphix Engine to leverage the security features offered by databases, operating systems, storage devices, and networks.

Anticipate external Attack Vectors: Combat attacks sourced from connected systems.

Enforce strong credentials: Define and enforce password policies.

Related Topics

- [Security](#)

Product Security

This section covers the following topics:

- [Delphix Product Security](#)
- [Software Updates](#)
- [Network Security](#)
- [Password Vault Support](#)
- [Certificate Management](#)
- [Replication Security](#)

Delphix Product Security

Overview

Delphix takes an active approach towards the discovery of vulnerabilities by employing third-party VAPT (vulnerability and penetration testing) Teams on an annual basis, as well as actively reviewing additional tools to assess our products. We also actively monitor bugs and vulnerabilities against the list of open-source software which is integrated into our products.

Delphix investigates all known vulnerabilities, either found by customers or internally, or discovered by the VAPT Teams. During the course of this investigative process, Delphix works with the reporter of the vulnerability to gather the technical information and determine the appropriate remedial action.

Customers can perform their own security scans and audits using their access through application administrative accounts. Because operating system administrative access is not provided with the closed software virtual appliance, scans executed with this privilege level are consistent with the level of access available in the environment.

As needed, we provide compensating controls and mitigations to reduce the impact of security issues more quickly and deliver fixes as part of the standard software release process described above.

Software Delivery Security

We deliver our virtual software appliance and upgrade images to customers on a secure server with access control. Additionally, we provide cryptographic signatures for each image, and customers can use these signatures to ensure the software images have not been tampered with.

Ancillary Components

The virtual appliance communicates with target servers running virtual databases (VDBs). The software is pushed to target servers to facilitate communication for Oracle and SQL Server databases. (Oracle toolkit and Delphix connector, respectively). This software lives outside of the DDDP product appliance, and new versions of the DDDP product will provide updated versions of this software as well.

Related Topics

- [Security](#)

Software Updates

Keeping Software up to date is an important part of any hardening plan. Delphix software releases are cumulative and include bug fixes, new features, and security improvements. In addition, Delphix releases hotfixes, procedures, and workarounds for critical vulnerabilities.

Patch Annually

To stay up to date, patch your Delphix Engine at least once per year. If you do not, you might have to upgrade twice to get the latest releases, and your old installed version will not be able to receive vulnerability fixes.

If possible, patch more frequently. Depending on the version you are upgrading from/to, you may be able to avoid or defer the reboot sequence, which defers downtime for your virtual databases (VDBs). This allows you to patch outside of downtime windows.

Subscribe to Delphix Notifications

Delphix issues email notifications when critical vulnerabilities are discovered. Registered support accounts will automatically receive these notifications. To ensure that you receive these notifications:

- Register at least two Engine Admins with Delphix Support
- Add Delphix Support accounts when Engine Admins leave the company

Related Topics

- [Security](#)

Network Security

Review the official documentation for the full list of required ports, *which depends on your database vendor*. Open only those ports that are required. The following table only lists generic requirements; you will need additional ports to integrate with databases.

Ports

Open Only Required Ports.

General Port Allocation

The Delphix Engine makes use of the following network ports irrespective of the type of database objects on it:

General Outbound Port Allocations

Protocol	Port Numbers	Use
TCP	21	Passive FTP connections from the Delphix Engine to the Delphix FTP server. Used for sending logs to Delphix Support.
TCP	25	Connection to a local SMTP server for sending email.
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
TCP/UDP	389	Standard access to an LDAP server
TCP/UDP	636	Secure access to an LDAP server
TCP	1023	Used to complete the passive FTP connection from the Delphix Engine to Delphix Support. These ports are not used if Delphix uses a proxy server to connect to the Internet.
TCP	8415	A Delphix replication source will connect to the replication target using this destination port

General Inbound Port Allocations

Protocol	Port Number	Use
TCP/UDP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
TCP	443	HTTPS connections to the Delphix GUI
TCP	8415	A Delphix replication target will accept incoming connections to this port from the replication

Related Topics

- [Security](#)

Password Vault Support

Overview

More and more organizations use Enterprise Password Vaults (EPV) such as CyberArk and HashiCorp Vault to store securely and centrally manage identities and credentials. Delphix has added CyberArk and HashiCorp Vault support to the Delphix Virtualization Engine as a new authentication option for environments and databases. This minimizes the number of places where credentials need to be stored and, therefore the risk of insecure storage.

The Delphix Engine uses various authentication methods such as username/password, username/ssh key, and Kerberos credentials when connecting to hosts and databases from the Engine. These credentials are stored on the Delphix Engine in an encrypted format and can be retrieved later to perform various operations. Delphix provides an additional authentication method by integrating Virtualization with the most common vault types (CyberArk, HashiCorp). At runtime, Delphix retrieves the credentials (passwords, ssh keys) from the customer's vault servers via API calls and avoids managing customer passwords.

Configuring Password Vaults

In the Setup app, system administrators can manage (add, delete, modify, and validate) vault configurations during and after the initial setup. Each engine can have multiple vaults configured of any type.

The authentication method supported for CyberArk is Certificate-based. Configuring a CyberArk vault requires providing a host address, port number, applicationId, and a client certificate (certificate chain and private key).

The authentication methods supported for HashiCorp vaults are Token-based, AppRole-based, and Certificate-based. Configuring a HashiCorp vault requires providing TokenId/ RoleId-SecretId/Certificates along with host address, port number, and namespace. Support for HashiCorp namespaces was added in Delphix 6.0.6.0.

Using Password Vaults

The Virtualization engine retrieves credentials at runtime from a vault using a unique identifier that locates a set of credentials in a configured vault. This occurs for any activity that requires Environment access (SnapSync, Validated Sync, LogSync, as well as Environment monitoring). This may result in a significant number of requests, so any existing connection rate limits should be evaluated and adjusted accordingly. For CyberArk, the unique identifier consists of a query string. For HashiCorp Vault, it consists of four parameters: engine, path, and a pair of keys that locate the username and secret (password or SSH key) in the key-value store at that engine and path.

To set up an environment or database user to use a vault, use the credential type **VaultCredential** when adding/modifying such users and specify the vault and the unique identifier of the credentials.

Roles and Privileges for CyberArk and HashiCorp Users

Role	Privileges
System Administrator	Can add, modify, delete, and list vault configurations.
Delphix Administrator	Can list existing vault configurations and link environment and database users to vault credentials.

Supported Environments and Databases


All environment users can use vault credentials. For Windows, the initial link via the Delphix Connector does not support vaults, but environment users can be subsequently updated to use vaults.

Vault integration is currently supported for SAP ASE database users, Oracle database users, and MSSQL domain users using 6.0.4 and later.

Setting up a Vault via GUI

Complete setup via the GUI is available for CyberArk as of 6.0.3.0 and for HashiCorp as of 6.0.4.0.

1. Connect to the Delphix Engine [http:// <Delphix Engine>/login/index.html#serverSetup](http://<Delphix Engine>/login/index.html#serverSetup).
2. Add a CyberArk or HashiCorp CA certificate to the TrustStore as part of the initial configuration. Refer [TrustStore Settings](#) for steps to add a CA certificate.
3. Click on the **Modify** link in the top right of the **Network Authorization** panel.



4. On the Network Authorization window, click "+" to add a new vault.
5. Enter the following information:
Depending upon your requirements, you can set the configurations for CyberArk or HashiCorp vault.

Field	Possible Value and Data Type	Description
Vault Type -CyberArk		

Field	Possible Value and Data Type	Description
Vaults Name	<user-specified> Accepts a string value	Specifies the user-specified vault name
Vault Hostname	cyberark02.delphix.com Accepts a URL string value	Specifies the location of the user's vault server
Port	443 Accepts an integer value	Specifies the port number through which the communication will happen
App ID	DelphixQADSA Accepts a string value	Specifies an application ID registered with and provided by CyberArk
Authentication Certificate	-----BEGIN CERTIFICATE----- <certificate> -----END CERTIFICATE----- Accepts a string value	Specifies the authentication certificate provided by CyberArk
Private Key	<CyberArk-provided> Accepts a string value	Specifies the private key provided by CyberArk for TLS based authentication
Vault-Type - HashiCorp		
Authentication Method	Token	
Vault Name	<user-specified> Accepts a string value	Specifies the user-specified vault name
Vault Hostname	10.119.132.39 Accepts a URL string value	Specifies the location of the customer vault server
Port	8100 Accepts an integer value	Specifies the port number through which the communication will happen
Vault Namespace	purple Accepts a string value	Specifies the namespace configuration specific to the user environment that is provided by the HashiCorp Enterprise Platform

Field	Possible Value and Data Type	Description
Token	s.2xvGgVRW53Ug6TXbCB7hj6a 6.waR7a Accepts a string value	Specifies the token specific to the user environment that is provided by the HashiCorp Enterprise Platform
Authentication Method	AppRole	
Vault Name	<user-specified> Accepts a string value	Specifies the user-specified vault name
Vault Hostname	10.119.132.39 Accepts a URL string value	Specifies the location of the customer vault server
Port	8100 Accepts an integer value	Specifies the port number through which the communication will happen
Vault Namespace	purple Accepts a string value	Specifies the namespace configuration specific to the user environment. This feature is provided with the HashiCorp Enterprise Platform
RoleID	f02b1a68-2c62-1230- e8dc-92424dea6468 Accepts a string value	Specifies the RoleID specific to the user environment
SecretID	ca3aca0c-254a-3f61- c943-0210fd670658 Accepts a string value	Specifies the SecretID specific to the user environment
Authentication Method	Certificate	
Vault Name	<user-specified> Accepts a string value	Specifies the user-specified vault name
Vault Hostname	10.119.132.39 Accepts a URL string value	Specifies the location of the customer vault server

Field	Possible Value and Data Type	Description
Port	8100 Accepts an integer value	Specifies the port number through which the communication will happen
Vault Namespace	purple Accepts a string value	Specifies the namespace configuration specific to the user environment. This feature is provided with the HashiCorp Enterprise Platform
Authentication Certificate	-----BEGIN CERTIFICATE----- <certificate> -----END CERTIFICATE----- Accepts a string value	Specifies the authentication certificate provided by HashiCorp for TLS based authentication
Private Key	ca3aca0c-254a-3f61- c943-0210fd670658 Accepts a string value	Specifies the private key specific to the user environment that is provided by HashiCorp for TLS based authentication
Role Name (Optional)	purple-admin-role Accepts a string value	Specifies the certificate role name for TLS based authentication

6. Click **Validate** to check the configurations before saving the vault details.
The below screenshot shows an example of the HashiCorp Vault configuration.

Add Password Vault ✕

Vault Type
HashiCorp ▾

Authentication Method
Token ▾

Vault Name

Vault Hostname

Port
▾


Vault Namespace

Token

Valid Vault

7. Click **Save**.

The added configurations can be viewed in the Network Authorization window.



Network Authorization

Kerberos Configuration
 Use Kerberos authentication to communicate with remote hosts

Host Connection Authentication
 When connecting to hosts, you can provide username-password pairs when setting up the connection, or you can utilize one or more Enterprise Password Vault systems by adding them to your engine setup.
 Click the + to add a vault

Vault name	Hostname	Port	Vault Type	Auth Method
TestVault	10.119.132.40	8200	HashiCorp	Token
TestVault1	10.119.132.40	8200	HashiCorp	AppRole
TestVault2	10.119.132.40	8200	HashiCorp	Certificate


Cancel Save

Editing a Vault via the GUI

1. Connect to the Delphix Engine [http:// <Delphix Engine>/login/index.html#serverSetup](http://<Delphix Engine>/login/index.html#serverSetup).
2. Click on the **Modify** link in the top right of the **Network Authorization** panel.
3. In the **Network Authorization** window, select a vault, then the **pencil** icon.
4. Edit your configuration.
5. Select **Edit**.

Deleting a Vault via the GUI

1. Connect to the Delphix Engine [http:// <Delphix Engine>/login/index.html#serverSetup](http://<Delphix Engine>/login/index.html#serverSetup).
2. Click on the **Modify** link in the top right of the **Network Authorization** panel.
3. In the **Network Authorization** window, select a vault, then select the **trashcan** icon.
4. Select **Yes** to delete the vault.



Delete Password Vault DemoVault

Are you sure you want to delete DemoVault ?

No Yes

Adding a Host user for HashiCorp

1. Login to the **Delphix Management Application** and select **Manage > Environments**.
2. Select Add **Environment**.
3. In the **Environment Setting** tab, select **Password Vault** as the Login Type.
4. Select the vault configuration and provide the secret engine name, path, and keys for the username secret and complete your environment configuration.
5. The environment will be created with the primary user using vault credentials.

The screenshot shows the 'Add Environment' dialog box with the following configuration:

- Host and Server:** Not selected
- Environment Settings:** Selected
- Summary:** Not selected
- Database Authentication:**
 - Database Username and Password
 - Username and Public Key
 - Password Vault
- Username:** sybaseUser
- Select the Enterprise Password Vault system:** TestVault
- Engine:** kv-v2/delphix-secrets
- Path:** bbdhcp-sybase
- Username Key:** user
- Secret Key:** pass

Buttons at the bottom: Cancel, Back, Next, Submit.

Adding a Database user for HashiCorp

1. Login to the **Delphix Management Application**.
2. Add dSource using database credentials from HashiCorp vault by selecting **Password Vault** as the **Login Type**.
3. Provide the appropriate secret engine name, path, and keys for the username secret and complete configuration.

Adding a Host user for CyberArk

1. Login to the **Delphix Management Application** and select **Manage > Environments**.
2. Select Add **Environment**.
3. In the **Environment Setting** tab, select **Password Vault** as the Login Type.
4. Select the vault configuration and provide the username, select the enterprise password vault system, and enter a query string that is a unique identifier pointing to the credentials to be retrieved and complete your environment configuration.

- The environment will be created with the primary user using vault credentials.

Add Environment

- Host and Server
- Environment Settings
- Summary

Host Address

SSH Port

Login Type

Username and Password
 Username and Public Key
 Password Vault

Username Ⓜ

Select the Enterprise Password Vault system

DemoVault ▼

Enter the Query String

Valid Credentials

Adding a Database user for CyberArk

1. Login to the **Delphix Management Application**.
2. Add dSource using database credentials from CyberArk vault by selecting **Password Vault** as the **Login Type**.
3. Provide the appropriate query string and complete configuration.

Add dSource

- Source
- dSource Configuration
- Data Management
- Policies
- Hooks
- Summary

SAP ASE 15.5 ESD#5.1

- loaddb
rhel-env
SAP ASE 15.7 SP138
- multi_mixed
rhel-env
SAP ASE 15.7 SP138
- one_data_two_log
rhel-env
SAP ASE 15.7 SP138

Instance
ASE1550_S5

Database
fuji

Page Size
4096

Environment User

anyname ▼

Login Type

Username and Password
 Password Vault

Username Ⓜ

Select the Enterprise Password Vault system

DemoVault ▼

Enter the Query String

Valid Credentials

Setting up a Vault via CLI

1. Login as a system administrator and add a CyberArk or HashiCorp CA certificate to the TrustStore as part of the initial configuration.

Network Security [Settings](#)

KeyStore	TrustStore	Open CSRs	...
D. Name	Issued ...	Expires	
CN=Engine ...	CN=Engine ...	Apr 20, 202...	
CN=cyberar...	CN=cybera...	Jan 10, 202...	

CN=cyberark.local, OU=BizDev, O=Cyberark, L=Newton, ST=MA, C=US [More info](#)

Issued To CN=cyberark.local, OU=BizDev, O=Cyberark, L=Newton, ST=MA, C=US

Issued By CN=cyberark.local, OU=BizDev, O=Cyberark, L=Newton, ST=MA, C=US

Serial Number 6076310049607000360

Valid From Jan 10, 2019 8:16:00 PM UTC

Expires Jan 10, 2029 8:16:00 PM UTC

[Certificate Information?](#)

2. SSH to `<Delphix Engine IP>service passwordVault` and enter **create**.

```
ip-10-110-230-197 service passwordVault>create
```

3. Add a new vault configuration by entering a name, host, port, applicationId, client certificate, and private key.

```
ip-10-110-230-197 service passwordVault create
Properties
type: CyberArkPasswordVault
```

```

name: DemoVault (*)
applications: Delphix (*)
clientCertificate:
  type: PemClientCertificate (*)
  clientCertificateChain: (required)
  privateKey:(required)
host: services-uscentral.sktyop.com (*)
port: 17993 (*)

```

4. Add a new HashiCorp vault configuration by entering a name, host, port, and other authentication information based on the authentication method (Token/AppRole/Certificate) selected.

HashiCorp - Token Based Authentication

```

ip-10-110-230-197 service passwordVault create
Properties
  type: HashiCorpVault (*)
  name: HashiCorpDemoVault (*)
  authentication:
    type: HashiCorpTokenAuthentication (*)
    token: ***** (*)
  host: 10.119.132.40 (*)
  port: 8200 (*)

```

HashiCorp - AppRole Based Authentication

```

ip-10-110-230-197 service passwordVault create
Properties
  type: HashiCorpVault (*)
  name: HashiCorpDemoVault (*)
  authentication:
    type: HashiCorpAppRoleAuthentication (*)
    roleId: 20d19a46-6fd9-c78b-b7e3-e43be4c8d5c2 (*)
    secretId: ***** (*)
  host: 10.119.132.40 (*)
  port: 8200 (*)

```

HashiCorp - Certificate Based Authentication

```

ip-10-110-230-197 service passwordVault create
Properties
  type: HashiCorpVault (*)
  name: HashiCorpDemoVault (*)
  authentication:
    type: HashiCorpCertificateAuthentication (*)
  clientCertificate:

```

```

    type: PemClientCertificate (*)
    clientCertificateChain: (required)
    privateKey: (required)
    roleName: (unset)
    host: 10.119.132.40 (*)
    port: 8200 (*)

```

Updating an existing vault configuration

```

ip-10-110-230-197 service passwordVault> select DemoVault
ip-10-110-230-197 service passwordVault 'DemoVault'>update
ip-10-110-230-197 service passwordVault 'DemoVault'update *> set name=TestVault
ip-10-110-230-197 service passwordVault 'DemoVault'update *> commit
ip-10-110-230-197 service passwordVault 'TestVault'>

```

Deleting an existing vault configuration

```

ip-10-110-230-197 service passwordVault 'TestVault'> delete
ip-10-110-230-197 service passwordVault 'TestVault' delete *> commit
ip-10-110-230-197 service passwordVault>

```

Adding/Modifying host users

Add an environment with user credentials from CyberArk vault. When adding a host/database user with a vault credential, the name field would be a user identifier and not the actual username. In case this field is empty, a unique identifier is generated with a hash of vault credentials.

```

ip-10-110-230-197 environment create *> set hostEnvironment.name=bbh-env
ip-10-110-230-197 environment create *> set hostParameters.host.address=bbdhcp-vault-
demo.dlpx.co
ip-10-110-230-197 environment create *> set hostParameters.host.toolkitPath="/work"
ip-10-110-230-197 environment create *> set primaryUser.name=oracleUser
ip-10-110-230-197 environment create *> set
primaryUser.credential.type=VaultCredential
ip-10-110-230-197 environment create *> set primaryUser.credential.vault=DemoVault
ip-10-110-230-197 environment create *> set
primaryUser.credential.vaultCredentialId="safe-test;folder=root;object=UnixSSH-
sybase"
ip-10-110-230-197 environment create *> commit
`UNIX_HOST_ENVIRONMENT -6
Dispatched job JOB-33
ENVIRONMENT_CREATE_AND_DISCOVER job started for "bbh-env".
ENVIRONMENT_CREATE_AND_DISCOVER job for "bbh-env" completed successfully.

```

Adding/Modifying **database users**

Add dSource using database credentials from CyberArk vault.

```

ip-10-110-230-197 database link *> set name=fuji
ip-10-110-230-197 database link *> set group=Untitled
ip-10-110-230-197 database link *> set linkData.config=ASE_SI_CONF-70
ip-10-110-230-197 database link *> set linkData.dbUser=sybaseUser
ip-10-110-230-197 database link *> set linkData.dbCredentials.type=VaultCredential
ip-10-110-230-197 database link *> set linkData.dbCredentials.vault=DemoVault
ip-10-110-230-197 database link *> set
linkData.dbCredentials.vaultCredentialId="safe-test;folder=root;object=Database-
Sybase-sa"
ip-10-110-230-197 database link *> set linkData.loadBackupPath='/opt/sybase/dumps"
ip-10-110-230-197 database link *> set linkData.sourceHostUser=HOST_USER-7
ip-10-110-230-197 database link *> set linkData.stagingHostUser=HOST_USER-7
ip-10-110-230-197 database link *> set linkData.stagingRepository=ASE_INSTANCE-6
ip-10-110-230-197 database link *> set
linkData.syncParameters.type=ASENewBackupSyncParameters
ip-10-110-230-197 database link *> commit
`ASE_DB_CONTAINER-1
Dispatched job JOB-39
DB_LINK job started for "Untitled/fuji".
DB_LINK job for "Untitled/fuji" completed successfully.

```

Update Existing Database Users

Convert an existing database to use vault credentials for the existing database user.

```

ip-10-110-230-197 > sourceconfig
ip-10-110-230-197 sourceconfig > select MyOraDB
ip-10-110-230-197 sourceconfig "MyOraDB" > update
ip-10-110-230-197 sourceconfig "MyOraDB" *> set
credentials.type=CyberarkVaultCredential
ip-10-110-230-197 sourceconfig "MyOraDB" *> set credentials.vault=MyVault
ip-10-110-230-197 sourceconfig "MyOraDB" *> set credentials.queryString="safe-
test;folder=root;object=UnixSSH-delphix_db"
ip-10-110-230-197 sourceconfig "MyOraDB" *> set db_user="Vault-User"
ip-10-110-230-197 sourceconfig "MyOraDB" *> commit

```

i The set db_user="Vault-User" is an optional step. If the db_user field is not changed, then it will continue to hold the old value. This value may no longer be correct, or the change to Vault credentials may represent an increase in the customer's security stance, and they may not want their Delphix Admins to know the username.

Setting up Vault via API

The vault API allows users to add, modify, delete, and list vault configurations and retrieving user credentials on a Delphix Engine.

Endpoint - <https://<Delphix Engine IP>/resources/json/delphix/service/passwordVault>

Sample API Request

```
{
  "type": "CyberArkPasswordVault",
  "name": "DemoVault",
  "host": "services-uscentral.skytap.com",
  "port": 17993,
  "applicationId": "Delphix",
  "clientCertificate": {
    "type": "PemClientCertificate",
    "privateKey": "-----BEGIN PRIVATE KEY-----<>-----END PRIVATE KEY-----",
    "clientCertificateChain": {
      "type": "PemCertificateChain",
      "chain": [
        {
          "type": "PemCertificate",
          "contents": "-----BEGIN CERTIFICATE-----<>-----END
CERTIFICATE-----"
        }
      ]
    }
  }
}
```

Deleting an existing vault configuration

```
{
  "type": "CyberArkPasswordVault",
  "name": "DemoVault",
  "host": "services-uscentral.skytap.com",
  "port": 17993,
  "applicationId": "Delphix",
  "clientCertificate": {
    "type": "PemClientCertificate",
    "privateKey": "-----BEGIN PRIVATE KEY-----<>-----END PRIVATE KEY-----",
    "clientCertificateChain": {
      "type": "PemCertificateChain",
      "chain": [
        {
          "type": "PemCertificate",
          "contents": "-----BEGIN CERTIFICATE-----<>-----END
CERTIFICATE-----"
        }
      ]
    }
  }
}
```

Adding/Modifying host users

Add an environment with user credentials from CyberArk vault. When adding a host/database user with a vault credential, the name field would be a user identifier and not the actual username. In case this field is empty, a unique identifier is generated with a hash of vault credentials.

```
{
  "type": "HostEnvironmentCreateParameters",
  "primaryUser": {
    "type": "EnvironmentUser",
    "credential": {
      "type": "VaultCredential",
      "vault": "CYBERARK_PASSWORD_VAULT-1",
      "vaultCredentialId": "safe=test;folder=root;object=UnixSSH-sybase"
    }
  },
  "hostEnvironment": {
    "type": "UnixHostEnvironment",
    "name": "bbh-env"
  },
  "hostParameters": {
    "type": "UnixHostCreateParameters",
    "host": {
      "type": "UnixHost",
      "address": "bbdhcp-vault-demo.dlpxdc.co",
      "toolkitPath": "/work"
    }
  }
}
```

Adding/Modifying **database users**

Add dSource using database credentials from CyberArk vault.

The following a sample API link request for MSSQL Domain User.

```
{
  "type": "LinkParameters",
  "name": "ReportServer",
  "group": "GROUP-1",
  "linkData": {
    "type": "MSSQLLinkData",
    "config": "MSSQL_SINGLE_CONFIG-5",
    "sharedBackupLocations": [],
  }
}
```

```

"encryptionKey": "",
"sourceHostUser": "HOST_USER-3",
"mssqlUser": {
  "password": {
    "type": "VaultCredential",
    "vault": "CYBERARK_PASSWORD_VAULT-2",
    "vaultCredentialId": "safe=test;folder=root;object=Database-MSSql-addtully"
  },
  "type": "MSSqlDomainUser"
},
"pptRepository": "MSSQL_INSTANCE-4",
"pptHostUser": "HOST_USER-3",
"ingestionStrategy": {
  "validatedSyncMode": "TRANSACTION_LOG",
  "type": "ExternalBackupIngestionStrategy"
},
"sourcingPolicy": {
  "logsyncEnabled": false,
  "type": "SourcingPolicy"
},
"syncParameters": {
  "compressionEnabled": false,
  "backupPolicy": "PRIMARY",
  "type": "MSSqlNewCopyOnlyFullBackupSyncParameters"
}
}
}

```

Certificate Management

The Delphix Dynamic Data Platform uses SSL certificates to secure a number of different communications. Delphix now provides users with the ability to manage their own certificates for HTTPS and DSP (Delphix Session Protocol) connections to and from the Delphix Engine. Users are able to see multiple certificates, their attributes, and deployed status, they can also:

- Replace certificates
- Delete certificates
- Create CSRs (certificate signing requests)

Configuring Network Security Settings

To enable and modify the extent of security settings for HTTPS and DSP (Delphix Session Protocol), click on the “Settings” button in the top right of the “Network Security” panel. This will open a new wizard where you can modify the settings.

Configuring settings for HTTPS and DSP.

The following procedure guides you through the process of enabling or modifying security settings for HTTPS and DSP.

1. Connect to the Delphix Engine <http://<Delphix Engine>/login/index.html#serverSetup>
2. Click on the **Settings** link in the top right of the **Network Security** panel. This will open a new wizard where you can modify the settings.

Network Security

[Settings](#)

KeyStore
TrustStore
Open CSRs
...

D. Name	Issued By	Expires
CN=Engine ip-10-11...	CN=Engine ip-10-11...	Sep 7, 2025 7:05:10 ...

CN=ENGINE IP-10-110-209-182.DELPHIX.COM CA, C=US [More info](#)

Is Accepted	Yes
Issued To	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Issued By	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Serial Number	94617914
Valid From	Sep 7, 2021 7:05:10 AM UTC
Expires	Sep 7, 2025 7:05:10 AM UTC

[Certificate Information](#)

CUSTOM AUTHORIZATIONS

Perform Server (target engine) authorization for Replication
Disabled

Perform Client (source engine) authorization for Replication
Disabled

Perform Server (this engine) authorization for remote connections
Disabled

Perform Client (target host) authorization for remote connections
Disabled

Perform Server (this engine) authorization for network test remote connections
Disabled

Perform Client (target engine or host) authorization for network test remote connections
Disabled

3. In the **Network Settings** wizard, you can modify the **HTTP mode**, **TLS Version**, and **TLS Cipher Suites** to be used.

The screenshot shows the 'Network Security Settings' wizard with the 'HTTP Configuration' step selected. The left sidebar contains three steps: 'HTTP Configuration' (selected), 'Custom Authorizations', and 'Summary'. The main content area is titled 'HTTP Configuration' and contains the following settings:

- HTTP Mode**: A dropdown menu set to 'HTTP and HTTPS'.
- TLS Version**: A dropdown menu set to 'TLSv1.2'.
- HTTPS Ciphers**: A checkbox labeled 'Select All' is unchecked, and a dropdown menu below it is set to 'TLS DHE RSA with AES 256 GCM SHA384, TLS DH...'.

4. Click **Next**.
5. By default, self-signed certificates are used for server authentication and username/password login acts as client authorization. You can enable the use of custom certificates and modify the extent of how they are used for the following features:
- [Replication](#)
 - [Remote Host Connections](#) (SnapSync, Remote hosts for Target or Server environments, and Oracle V2P)
 - [Throughput tests](#) using the Network Performance Tool and DSP

 For all three features, Server authentication must be enabled for Client authentication to take effect.

 Any configuration change requires a stack restart to take effect. Submitting the Network Security Settings in the GUI will automatically trigger a stack restart. All jobs will be stopped, but VDBs will continue to run.

Network Security Settings

- HTTP Configuration
- Custom Authorizations**
- Summary

Custom Authorizations

Use your own certificates to authorize DSP connections for Server and/or Client.
Follow the process described [here](#)

Replication ⓘ

- Perform Server (target engine) authorization for Replication
- Perform Client (source engine) authorization for Replication

Remote Connections (SnapSync, Remote Hosts, Oracle V2P) ⓘ

- Perform Server (this engine) authorization for remote connections
- Perform Client (target host) authorization for remote connections

Throughput Tests (Network Performance Tool) ⓘ

- Perform Server (this engine) authorization for network test remote connections
- Perform Client (target engine or host) authorization for network test remote connections

6. Once the HTTP and DSP configurations have been set as desired, click **Next**.
7. You will be presented with a Summary tab. Clicking **Submit** to accept your changes will trigger a stack restart as this is necessary for the configuration changes to take effect. Note: all jobs will be stopped, but VDBs will continue to run.

Related Topics

- [Configuring Network Security Settings](#)

Certificate Management and Remote Connections

Overview

The Server is the engine and the Client is the Remote Host. The protocol can be used for SnapSync, Oracle V2P (Virtual to Physical), and remote host connections. Once either of these options is enabled, the steps for setting up the remote keystore/truststore must be done for any future environments added to the engine.

Enabling Server Authentication

To enable Server Authentication, do the following:

1. Replace the desired certificate for DSP (Delphix Session Protocol) in the engine KeyStore. For more details, refer to [KeyStore Settings](#).
2. Create a JKS or PKCS#12 keystore on the remote host with the full CA chain of the replaced certificate. Make sure the created keystore has permissions such that it is readable by all environment users, then enter the keystore details into the host's truststore configuration on the engine. For more details, refer to [Host DSP Configuration](#).
3. Select the option **Perform Server (this engine) authorization for remote connections**.

Enabling Client Authentication

To enable Client Authentication, enable Server Authentication (refer to above steps), then do the following:

1. Create a JKS or PKCS#12 keystore on the remote host with the desired key pair. Make sure the created keystore has permissions such that it is readable by all environment users, then enter the keystore details into the host's keystore configuration on the engine. For more details, refer to [Host DSP Configuration](#).
2. Add the full CA chain of the remote host's key pair to the TrustStore on the engine. For more details, refer to [TrustStore Settings](#).
3. Select the option **Perform Client (the target host) authorization for remote connections**.
4. Once the configurations have been set as desired, you will be presented with a summary page. Clicking **Submit** will trigger a stack restart as that is necessary for the configuration changes to take effect. Note: all jobs will be stopped, but VDBs will continue to run.

Related Topics

- [Configuring Network Security Settings](#)
- [Host DSP Configuration](#)
- [TrustStore Settings](#).
- [KeyStore Settings](#)

Certificate Management and Replication

In Replication, the Server is the Target engine and the Client is the Source engine.

Enabling Server Authentication

To enable Server Authentication, do the following:

1. Replace the desired certificate for DSP (Delphix Session Protocol) in the Target engine KeyStore. For more details, refer to [KeyStore Settings](#).
2. Add the full CA chain of the replaced certificate from the Target engine to the TrustStore on the Source engine. The CA chain must match on both engines. For more details, refer to [TrustStore Settings](#).
3. Select the option **Perform Server (target engine) authorization for Replication** for both Target and Source engines.

Enabling Client Authentication

To enable Client Authentication, enable Server Authentication (refer to above steps), then do the following:

1. Replace the desired certificate for DSP in the Source engine KeyStore. For more details, refer to [KeyStore Settings](#).
2. Add the full CA chain of the replaced certificate from the Source engine to the TrustStore on the Target engine. The CA chain must match on both engines. For more details, refer to [TrustStore Settings](#).
3. Select the option Perform Client (source engine) authorization for Replication for both Target and Source engines.
4. Once the configurations have been set as desired, you will be presented with a summary page. Clicking **Submit** will trigger a stack restart as that is necessary for the configuration changes to take effect. Note: all jobs will be stopped, but VDBs will continue to run.

Related Topics

- [Configuring Network Security Settings](#)
- [TrustStore Settings](#)
- [KeyStore Settings](#)

Certificate Management for Throughput Tests

Enabling Server Authentication

To enable Server Authentication, do the following:

1. If using the engine to engine tests, follow the steps for [Replication Server Authentication](#).
2. If using the engine to host tests, follow the steps for [Remote Connections Server Authentication](#).
3. Select the option **Perform Server (this engine) authorization for remote connections**.

Enabling Client Authentication

To enable Client Authentication, enable Server Authentication (refer to above steps), then do the following:

1. If using the engine to engine tests, follow the steps for [Replication Client Authentication](#).
2. If using the engine to host tests, follow the steps for [Remote Connections Client Authentication](#).
3. Select the option **Perform Client (target engine or host) authorization for remote connections**.
4. Once the configurations have been set as desired, you will be presented with a summary page.
Clicking **Submit** will trigger a stack restart as that is necessary for the configuration changes to take effect.
Note: all jobs will be stopped, but VDBs will continue to run.


Related Topics

- [Replication Client Authentication](#)
- [Remote Connections Client Authentication](#)


Configuring HTTP settings for the Delphix Engine

Use the following steps to configure HTTP settings for the Delphix Engine.

1. Login to the Delphix Engine Setup UI as a sysadmin.
2. From the Dashboard under **Network Security** select **Settings**. This will open a new wizard where you can modify the settings.



3. In the **Network Security Settings** screen, you can modify the **HTTP mode**, **TLS Version**, and **HTTPS Ciphers** to be used.
4. In **HTTP mode**, select one of the following:
 1. **HTTP Only**: accepts only HTTP connections.
 2. **HTTPS Only**: accepts only HTTPS connections.
 3. **HTTP Redirect**: Redirect all requests made over HTTP to HTTPS.
 4. **HTTP and HTTPS**: accepts both HTTP and HTTPS connections.
 5. **HTTP Redirect with HSTS**: redirect all requests made over HTTP to HTTPS and add Strict-Transport-Security Header to all responses.



5. In **TLS Version**, select the required TLS versions.
6. In **HTTPS Ciphers**, select the required option from the drop-down list. Select the check box next to **Select All** if you want to select all the options.
7. Click **Next** and then click **Submit** to allow your engine to restart for the changes to take effect.
8. When the Delphix Engine homepage is loaded after the restart, all the HTTP requests will be redirected to the selected HTTP mode after the first load.

 The above steps might not work if:

1. The certificates are not trusted - 'not secure' warning on the search bar of the browser. Certificate setup should be completed before setting up HSTS headers for them to be useful - otherwise, the redirect will still happen over 302 status.
2. Your site is not prepared for HTTPS setup - i.e., each of the subdomains accessed should be supporting HTTPS - parts of the application might be inaccessible otherwise.
3. Every first request whenever the cache is cleared would redirect over 302. Any further requests after the first would redirect over 307.

Don't want to use HSTS?

1. By default engine is set up to 'HTTP and HTTPS'. You can leave the engine as it is and not make any changes.
2. If you set up HSTS and want to revert it, the same steps for setup can be followed from above, and select your previous option on step 3.

Related Topics

- [Configuring Network Security Settings](#)

Regenerating Self-Signed End-Entity and CA Certificates

 In many environments, the replacement of HTTPS and/or DSP may be unnecessary.

DSP certificate is only relevant if Custom Authorizations have been configured in Network Security settings, as discussed in the [Configuring Network Security Settings](#) article. If these checkboxes are not applied, this means the DSP certificate is not being used.

HTTPS certificate replacement is only necessary if HTTPS connections are used for web browser access.

The following process will leverage Java *keytool* utility. This is commonly available in most Java JDK installations, including those installed in the Delphix Toolkit for Unix, Linux, and Windows Environments under `<toolkit directory>/*host/java/jdk/bin/`. In the following example, `/work` is the toolkit directory. The subdirectory naming conventions from 5.3.x and 6.0.x are illustrated as:

```
$ find /work -name keytool
/work/Delphix_COMMON_f126df603015_33e2f61712c3_2_host/java/jdk/jre/bin/keytool
/work/Delphix_COMMON_f126df603015_33e2f61712c3_2_host/java/jdk/bin/keytool
/work/Delphix_COMMON_564d56b0_26ad_e6ac_f782_d15213207664_oracle_host/java/jdk/bin/
keytool
/work/Delphix_COMMON_564d56b0_26ad_e6ac_f782_d15213207664_oracle_host/java/jdk/jre/
bin/keytool
```

By the end of this process, a PKCS#12 keystore file is generated containing the CA certificate, DSP, and HTTPS certificate. This file will be used for upload twice in the System Setup interface.

Other notes:

- The recommended keystore password 'changeit' is used.
- For `<domain>`, replace this string with the Engine FQDN in every command. This is used as the CN (Common Name). For instance, `"-dname 'CN=Engine <domain> CA, C=US'"` would be replaced with `"-dname 'CN=Engine example.delphix.com CA, C=US'"`.
- The certificate aliases to be used are 'tomcat' for HTTPS and 'dsp' for DSP.
- The existing Delphix CA certificate in the truststore cannot be removed. Faults related to this certificate should be ignored.

1. Generate a new Delphix CA Certificate.

```
PASSWORD_ENV='changeit' keytool -genkeypair -noprompt -alias delphixca
-keyalg RSA -keysize 2048 -validity 397 -ext 1.3.6.1.5.5.7.3.1 -ext bc=ca:true
-ext ku=kCS,cRLS -sigalg SHA256withRSA -storepass:env PASSWORD_ENV -storetype
pkcs12 -startdate -10000M -dname 'CN=Engine <domain> CA, C=US' -keypass:env
PASSWORD_ENV -storetype pkcs12 -keystore keystore
```

2. Generate the HTTPS/TLS certificate

```
PASSWORD_ENV='changeit' keytool -genkeypair -alias tomcat -keyalg RSA -keysize
2048 -validity 397 -ext 1.3.6.1.5.5.7.3.1 -ext san=dns:<domain> -ext bc=ca:fa
lse -sigalg SHA256withRSA -storetype pkcs12 -storepass:env PASSWORD_ENV
```

```
-startdate -10080M -dname 'CN=<domain>, C=US' -keypass:env PASSWORD_ENV
-keystore keystore
```

3. Generate the DSP certificate

```
PASSWORD_ENV='changeit' keytool -genkeypair -alias dsp -keyalg RSA -keysize
2048 -validity 397 -ext 1.3.6.1.5.5.7.3.1 -ext san=dns:<domain> -ext bc=ca:fales
e -sigalg SHA256withRSA -storetype pkcs12 -storepass:env PASSWORD_ENV
-startdate -10080M -dname 'CN=<domain>, C=US' -keypass:env PASSWORD_ENV
-keystore keystore
```

4. Sign the HTTP/TLS certificate

```
PASSWORD_ENV='changeit' keytool -certreq -alias tomcat -keyalg RSA -sigalg
SHA256withRSA -storetype pkcs12 -keypass:env PASSWORD_ENV -storepass:env
PASSWORD_ENV -keystore keystore | PASSWORD_ENV='changeit' keytool -gencert
-alias delphixca -ext 'san=dns:<domain>' -validity 397 -sigalg SHA256withRSA
-storetype pkcs12 -storepass:env PASSWORD_ENV -keystore keystore -startdate
-10080M | PASSWORD_ENV='changeit' keytool -importcert -alias tomcat -storetype
pkcs12 -storepass:env PASSWORD_ENV -keystore keystore
```

5. Sign the DSP Certificate

```
PASSWORD_ENV='changeit' keytool -certreq -alias dsp -keyalg RSA -sigalg
SHA256withRSA -storetype pkcs12 -keypass:env PASSWORD_ENV -storepass:env
PASSWORD_ENV -keystore keystore | PASSWORD_ENV='changeit' keytool -gencert
-alias delphixca -ext 'san=dns:<domain>' -validity 397 -sigalg SHA256withRSA
-storetype pkcs12 -storepass:env PASSWORD_ENV -keystore keystore -startdate
-10080M | PASSWORD_ENV='changeit' keytool -importcert -alias dsp -storetype
pkcs12 -storepass:env PASSWORD_ENV -keystore keystore
```

At this point, the certificates can be installed by following the Customer Provided Key Pair method described in the [Customer Provided Key Pair Configuration](#) article.

The same keystore generated will be uploaded twice, once using alias 'dsp' and once using alias 'tomcat', to replace DSP and HTTPS certificates, respectively.

i If this error occurs, **Failed to read file with error "Invalid keystore format"**, ensure the **Upload certificate from a PKCS#12 keystore** radio button is selected.

Windows Keytool Distinctions

- Given the Delphix Connector installation directory **C:\Program Files\Delphix\DelphixConnector**, the keytool.exe executable can be found at **C:\Program Files\Delphix\DelphixConnector\jre\bin\keytool.exe**.
- Powershell set environment variable with:
 - \$ENV:PASSWORD_ENV='changeit'**
- Similar to the comment above with path assumption, on Windows, change ./keytool to .\keytool.exe if located in the jre\bin subdirectory.

Configuring Network Settings

This topic explains how to replace the HTTPS (HTTP Secure) certificate and the DSP (Delphix Session Protocol) certificate used by the Delphix Virtualization Engine. The KeyStore tab in the Network Security panel displays the keys that the Delphix Engine uses for its identity in HTTPS and DSP communication. For a selected service, the tab will show the full certificate chain starting from the end-entity certificate and ending with the root CA.

Network Security
[Settings](#)

KeyStore
TrustStore
Open CSRs
...

D. Name	Issued By	Expires
CN=Engine ip-10-11...	CN=Engine ip-10-11...	Sep 7, 2025 7:05:10 ...

CN=ENGINE IP-10-110-209-182.DELPHIX.COM CA, C=US [More info](#)

Is Accepted	Yes
Issued To	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Issued By	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Serial Number	94617914
Valid From	Sep 7, 2021 7:05:10 AM UTC
Expires	Sep 7, 2025 7:05:10 AM UTC

[Certificate Information](#)

CUSTOM AUTHORIZATIONS

- Perform Server (target engine) authorization for Replication**
Disabled
- Perform Client (source engine) authorization for Replication**
Disabled
- Perform Server (this engine) authorization for remote connections**
Disabled
- Perform Client (target host) authorization for remote connections**
Disabled
- Perform Server (this engine) authorization for network test remote connections**
Disabled
- Perform Client (target engine or host) authorization for network test remote connections**
Disabled

There are two methods of replacing a certificate. The key difference between the two is whether Delphix or the user is providing the key pair (public and private key).

For more information on how to configure these see:

- [Delphix Provided Key Pair Configuration](#)
- [User-Provided Key Pair Configuration](#)

Related Topics


- [Delphix Provided Key Pair Configuration](#)
- [User-Provided Key Pair Configuration](#)
- [Configuring Network Settings](#)

Customer Provided Key Pair Configuration

This section describes the steps to take if you are replacing the HTTPS or DSP (Delphix Session Protocol) with your own key pair and certificate. To start, you need to add the key pair and full certificate chain as an entry in a file in JKS or PKCS #12 format.

⚠ If replacing the key pair for the DSP application, note that when the DSP Server or Client authenticates each incoming DSP connection (if enabled), it will validate that each certificate of the incoming connection's identity chain has a "Valid From (or Not Before)" date that is after its own time. Thus, if your Delphix Engine or host environments are running off of an incorrect (slow) time configuration, then your DSP connections will not work until the offending engine or host's time advances past all incoming certificate's "Valid From (or Not Before)" time. If correcting the Delphix Engine's or host environment's time configuration may cause issues, then you can work around this issue by creating and using a certificate with a "Valid From (or Not Before)" date which is before your slowest Engine or host.

1. From the **Network Security** panel select the **Actions (...)** menu and select **Replace Certificate**.



2. In the **Replace Certificate** window, select **Upload certificate as files**.

Replace Certificate ✕

- Upload Certificate
- Summary

Upload Certificate

This will replace the HTTPS CN=js532.dcenter.delphix.com, C=US certificate on this Delphix Engine.

Note: only the X.509 standard, and JKS or PKCS#12 file formats are supported. Upon certificate replacement, the management stack will need to be restarted before the new certificate will be used.

You can either upload a keystore file containing the full trust chain, or you can paste the equivalent PEM file contents.

Upload certificate from a JKS keystore

Upload certificate from a PKCS#12 keystore

PEM file contents

Select the file to upload. Upload a single keystore file containing the chain of trust.

Certificate Alias

Keystore Passcode

Key Pair Passcode (if different from Keystore Passcode)

Click to browse
– OR –
Drag and drop a file here

Upload Status
File not selected

Note: The file will upload after you click Next.

Cancel Back Next Submit

- a. The **Certificate Alias** field is where the key pair and certificate is saved in your JKS or PKCS #12 store.
 - b. The **Keystore Passcode** field is the keystore’s password.
 - c. The **Key Pair Passcode** field is the password for the given alias’ key. If not set, it uses the keystore's password.
3. Click **Next** to view a summary. Then click **Submit**.


Related Topics

- [Configuring Network Settings](#)

Delphix Provided Key Pair Configuration

Use the following instructions to provide an HTTPS or DSP (Delphix Session Protocol) certificate chain for a key pair created by the Delphix Engine.

1. Connect to the Delphix Engine `HTTP://< your engine>/login/index/html#serverSetup`.
2. From the **Network Security** panel select the **Actions (...)** menu and select **Create CSR**.



3. In the Create CSR window provide options for your CSR (certificate signing request). The only required field is the "Common Name (CN)".

Create CSR ✕

Complete the following form and click "Create". A key pair and the CSR will be generated.
Once created, it can be accessed via the CLI.

Service

HTTPS
 DSP

Distinguished Name

Complete form fields
 Create a composite Distinguished Name

Common Name

Organizational Unit

Organization

City/Locality

State/Region

Country

- a. Select the service you want to create a CSR for (HTTPS or DSP)
- b. Select how to enter the distinguished name (using the provided fields, or as a composite string).

4. The Validate button will verify that your provided distinguished name follows a legal format.

Create CSR

Complete the following form and click "Create". A key pair and the CSR will be generated.
Once created, it can be accessed via the CLI.

Service

HTTPS

DSP

Distinguished Name

Complete form fields

Create a composite Distinguished Name

Composite Distinguished Name

test

Validate

Show advanced

Cancel Create

5. If desired, select the **Show advanced** link to provide extra options.
- Subject Alternative Names (SAN):** Add IP addresses or hostnames to be associated with the certificate. Some browsers require that the HTTPS certificate includes the server's fully qualified domain name (FQDN) as a SAN.
 - Force Replace:** By default, this is false and means Delphix will not replace the active key pair and certificate with the newly generated key pair and self-signed certificate. If the user wants to replace the active key pair right away before the signed certificate has been created this can be set to true.
 - Keypair Algorithm:** Choose whether the created keypair will use the RSA or ECDSA algorithm. Once an algorithm has been chosen the user can customize the key size and the signature algorithm used.
 - Key size:** The valid sizes for RSA are between 2048 and 4096 inclusive. The valid sizes for ECDSA are between 256 and 571 inclusive.
 - Signature algorithm:** The available signature algorithms for RSA are "SHA256withRSA", "SHA384withRSA", "SHA512withRSA". The available signature algorithms for ECDSA are "SHA256withECDSA", "SHA384withECDSA", "SHA512withECDSA".
6. Click **Create**.
- Once submitted, the CSR is created, and will show up in the **Open CSRs** tab. You can select the CSR and click on View PEM to obtain the CSR in a PEM format.

⚠ If creating a CSR for the DSP application, note that when the DSP Server or Client authenticates each incoming DSP connection (if enabled), it will validate that each certificate of the incoming connection's identity chain has a "Valid From (or Not Before)" date that is after its own time.


Thus, if your Delphix Engine or host environments are running off of an incorrect (slow) time configuration, then your DSP connections will not work until the offending engine or host's time advances past all incoming certificate's "Valid From (or Not Before)" time.

If correcting the Delphix Engine's or host environment's time configuration may cause issues, then you can workaround this issue by signing the CSR with a "Valid From (or Not Before)" date which is before your slowest Engine or host.

Replacing a Certificate

After the CSR has been signed and turned into an X.509 Certificate, you can replace the certificate using the Replace Certificate option

1. From the **Network Security** panel select the **Actions (...)** menu and select **Replace Certificate**.



2. In the Replace Certificate window, select PEM file contents. You can then paste the PEM response in the text box.

⚠ The PEM contents must contain a list of the entire trust chain from the newly generated end-entity certificate to the root CA. You can do this by just pasting each certificate back to back in the text box as long as they are separated by the begin and end cert tags. The order that the PEM certificates are added to the list does not matter.

3. Click **Next**.
4. The Summary tab provides a description of your selections. Click **Submit**.

Related Topics

- [Configuring Network Settings](#)

Host DSP Configuration

In Delphix Engine to remote host DSP (Delphix Session Protocol) communication, the Server is the engine and the Client is the host. The Add Environment wizard allows you to tell the Delphix Engine how the KeyStores and TrustStores have been set up on the remote hosts.

⚠ When the Server or Client authenticates each incoming DSP connection (if enabled), it will validate that each certificate of the incoming connection's identity chain has a "Valid From (or Not Before)" date that is after its own time.

Thus, if your Delphix Engine or host environments are running off of an incorrect (slow) time configuration, then your DSP connections will not work until the offending engine or host's time advances past all incoming certificate's "Valid From (or Not Before)" time.

If correcting the Delphix Engine's or host environment's time configuration may cause issues, then you can work around this issue by creating and using certificates with a "Valid From (or Not Before)" date which is before your slowest Engine or host.

Adding a new single instance environment after DSP configuration changes

- If server authentication for remote host communication or engine to host throughput tests is desired, make sure the appropriate config is set. For more details refer to [Configuring Network Security Settings](#). You will need to create a JKS or PKCS#12 keystore on the remote host with the full CA chain of the DSP key in the keystore. By default, the key will just be signed by the Delphix CA, but you can replace the DSP key if you wish. Refer to [KeyStore Settings](#) for more details.
- If client authentication for remote host communication or engine to host throughput tests is also desired, make sure the appropriate config is set. For more details refer to [Configuring Network Security Settings](#). You will need to create another JKS or PKCS#12 keystore on the remote host with the desired key pair. Make sure the created keystore has permissions such that it is readable by all environment users. Then, add the full CA chain of the remote host's key pair to the TrustStore on the engine. For more details, refer to [TrustStore Settings](#).

Once the appropriate toggles are enabled, and the remote host is all set up, you can now add the environment as shown below. If only server authentication was desired, only the TrustStore fields need to be filled in. If client authentication was also desired, then the KeyStore fields will also need to be filled in

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Actions (...)** menu next to Environments and select **Add Environment**.
5. In the **Host and Server** tab, select **Unix/Linux**.
6. Select **Standalone Host** or **Oracle Cluster**, depending on the type of environment you are adding.
7. Click **Next**.
8. In the Environment Settings tab enter your DSP configurations.
9. Select **Submit**.

Modifying an existing single instance environment after DSP configuration changes

1. If an environment already exists after enabling server/client DSP authentication, you will need to modify its attributes for host communication to continue working. As detailed in the above section [Adding a new single instance environment after DSP configuration changes](#), you will need to set up the appropriate stores on the remote host.
2. Once this is done, the **Details** page of your environment will show the existing DSP attributes.

3. Click the **pencil** icon in the top right corner to edit the DSP KeyStore and TrustStore fields accordingly.

Attributes 
Host Address bbdhcp-AHCI-58503.dcenter.delphix.com
NFS Addresses None
SSH Port 22
DSP KeyStore Path /work/my_keystore
DSP KeyStore Alias dsp
DSP TrustStore Path /work/my_truststore

4. Edit the DSP fields as required.
5. Once you are done select the checkmark.

Attributes

Host Address

NFS Addresses

SSH Port

DSP KeyStore Path ⓘ

DSP KeyStore Password ⓘ

DSP KeyStore Alias ⓘ

DSP TrustStore Path ⓘ

DSP TrustStore Password ⓘ

⚠ This will need to be done for all existing environments after DSP server/client authentication is enabled otherwise many host communication features will not work.

Adding a new Unix/Linux cluster environment after DSP configuration changes

Adding a new cluster environment is similar to adding a single instance environment, see [Adding a new single instance environment after DSP configuration changes](#). The steps for setting up the TrustStore and/or KeyStore will need to be done on ALL nodes of the cluster. For a new cluster, each node must also be set up to have the exact same path, password, an alias, because the Delphix Engine will use the same configuration for every auto discovered node. If desired, the path, password, and alias configuration can be changed for each node, but only AFTER the cluster has been added. See [Modifying an existing Unix/Linux cluster environment after DSP configuration changes](#) for more details.

Modifying an existing Unix/Linux cluster environment after DSP configuration changes

Similar to single instance environments, an existing cluster can be modified if DSP server/client authentication is enabled. In the cluster environment's Details page, each node can be selected and modified individually. You can also use this to change the path, password, and/or alias to be different across nodes for a cluster that was just added.

 This will need to be done for all existing environments after DSP server/client authentication is enabled otherwise many host communication features will not work.

DSP for Windows Clusters

Before adding a Windows cluster, each node must have already been added as a single instance. See [Adding a new single instance environment after DSP configuration changes](#) if adding a new Windows cluster. If modifying a Windows cluster after a DSP configuration change, you can modify each Windows node on its environment details page. See [Modifying an existing single instance](#) for more details, or just use the cluster environment details which will allow you to select which node to modify. See [Modifying an existing Unix/Linux cluster environment after DSP configuration changes](#).

KeyStore Settings

The KeyStore tab in the Network Security panel displays the public key certificates that the Delphix Engine uses for its identity in HTTPS and DSP (Delphix Session Protocol) communication. For a selected service, the tab will show the full certificate chain starting from the end-entity certificate and ending with the root CA.

Network Security
[Settings](#)

KeyStore
TrustStore
Open CSRs
...

D. Name	Issued By	Expires
CN=Engine ip-10-11...	CN=Engine ip-10-11...	Sep 7, 2025 7:05:10 ...

CN=ENGINE IP-10-110-209-182.DELPHIX.COM CA, C=US [More info](#)

Is Accepted	Yes
Issued To	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Issued By	CN=Engine ip-10-110-209-182.delphix.com CA, C=US
Serial Number	94617914
Valid From	Sep 7, 2021 7:05:10 AM UTC
Expires	Sep 7, 2025 7:05:10 AM UTC

[Certificate Information](#)

CUSTOM AUTHORIZATIONS

- Perform Server (target engine) authorization for Replication**
Disabled
- Perform Client (source engine) authorization for Replication**
Disabled
- Perform Server (this engine) authorization for remote connections**
Disabled
- Perform Client (target host) authorization for remote connections**
Disabled
- Perform Server (this engine) authorization for network test remote connections**
Disabled
- Perform Client (target engine or host) authorization for network test remote connections**
Disabled


There are two methods of replacing a certificate. The key difference between the two is whether Delphix or the user is providing the key pair (public and private key). For more information see [Configuring Network Settings](#).

Related Topics

- [Certificate Management](#)


TrustStore Settings

The TrustStore tab in the Network Security panel displays all the CA certificates that the Delphix Engine trusts. Click on the Actions (...) menu in the top right to reveal the options for editing the TrustStore contents:



Adding a Certificate

1. From the Actions menu select Add Certificate.
2. In the Add Certificate wizard paste the PEM contents of the CA certificate you want to add. The PEM contents must have the appropriate header and footer included.



3. If you are adding a non-root CA certificate, its signer must already exist in the truststore. So, if you are adding a chain with multiple certificates, you must add them individually starting from the root CA. If not, you will get an error saying that we could not establish a chain of trust.

- Click **Next** to view a Summary tab where you can confirm the certificate contents.

Add Certificate

Upload Certificate

Summary

Summary

NEW CERTIFICATE

Is Accepted

Issued To

Issued By

Serial Number

Valid From

Expires

SHA1 Fingerprint

MD5 Fingerprint

Is Certificate Authority

CN=

CN=

0

Jul 15

Aug 14

Yes

Cancel Back Next Submit

- Click **Submit**.

Deleting a Certificate

Use this option to delete the selected CA certificate.

⚠ Deleting any certificate that breaks an existing chain of trust is not allowed.

- From the Actions menu select **Delete Certificate**.
- The default Delphix CA cannot be deleted. Note that deleting any certificate that breaks an existing chain of trust is also not allowed.
- In the Confirmation dialog select **Delete**.

Delete Certificate

Are you sure you want to delete certificate CN=Engine js532.dcenter.delphix.com CA, C=US?

Cancel Delete

⚠ Unlike the network security settings, any changes to the TrustStore will not require a stack restart.

Related Topics

- [Certificate Management](#)

Replication Security

Choose Encrypt and Compress when Replicating

Delphix provides encryption capabilities when replicating data from one Delphix Engine to another.

Upon selecting the **Encrypted** option under **Traffic Options** the Delphix Engine intelligently compresses before encrypting data. This ordering leads to lower CPU utilization and higher throughput compared to using encryption alone, with the same level of protection. See [Configuring Replication](#) for details. The flag can be set at any time. Once set, it results in sending subsequent replication traffic streams as encrypted.

Train your Delphix Administrators to choose both options.

Related Topics

- [Configuring Replication](#)
- [Security](#)

Object Security

This section covers the following topics:

- [User Management](#)
- [Source Database Security](#)
- [Source and Target Host Security](#)

User Management

Secure User Management is best achieved by integration with your centralized authentication service. Once the integration is complete, create LDAP authenticated named users to facilitate separation of duties, least privileges, and auditing. Disable the out-of-the-box generic ADMIN and SYSADMIN accounts.

Use LDAP for Authentication

As described under System Configuration above, enable LDAP authentication to leverage your enterprise authentication service and enable SSL/TLS to secure LDAP connections.

Create Named Users

Do not create generic functional accounts such as “QA,” “DEV,” or “TEST.” Such accounts will not leave a proper audit trail and violate the separation of duties principle. Instead, create LDAP authenticated named users.

Assign Least Privileges

Restrict the **admin** and **sysadmin** roles to 1-2 trusted named users each. These roles are highly privileged and must be carefully managed. These roles typically map to a **DBA** and **System Administrator** respectively.

For subordinate users who need to refresh VDBs, assign “Data Operator” privilege on the VDB and “Reader” privilege on the dSource.

For subordinate users who need to provision new VDBs from dSources, assign “Provisioner” privilege on the dSource and “Provisioner” privilege on the Group to which they will assign the VDB.

Consider Delphix Self-Service Functionality

The Delphix Self-Service functionality is targeted towards developer and tester self-service, and it contains a more sophisticated privilege model. With this functionality, Delphix Self-Service users do not have access to the Administrator GUI.

Administrators can define multiple data sources as a complete template. They also allocate server resources as a “data container.” The end-user has the ability to update data from the source, from peers using the same source, and from prior images of the source that they have created.

Disable ADMIN and SYSADMIN

Once you have established named Delphix Administrators and Systems Administrators, disable the out-of-the-box `sysadmin` and `admin` accounts. You can disable accounts through the CLI.

 When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

Related Topics

- [Object Security](#)

Source Database Security

Choose Minimum Privileges for Delphix DB User

The Delphix Engine requires a DB user (**delphix_db**) on your source databases. This account is necessary to detect the state of the source and stay in sync. Do not give unnecessary privileges to this user. Leverage the script provided by Delphix in the hostchecker bundle to create a user with the minimum required privileges. The DB user (e.g., **delphix_db**, which is the example used on this page) account can have the same or different user name on each of your source databases.

Protect the Delphix DB User Password

Because the **delphix_db** user has access to sensitive data dictionary information, take steps to protect access to this account.

Use Database Encryption Functionality to Encrypt Sensitive Data at Rest

Database vendors provide tools to encrypt data at rest. This encrypts data on disk in order to provide protection at rest. Use the database encryption on your source databases to encrypt sensitive or all data. Delphix integrates seamlessly with database encryption to provide protection at rest.

Choose Encrypt and Compress when Linking

Delphix provides encryption capabilities when linking against your source databases. Encrypting while linking can lead to higher CPU utilization and higher throughput. This overhead can be reduced by selecting both the Compress and Encrypt options, in order to compress the data before it is encrypted. Train your Delphix Administrators to choose both options.

Related Topics

- [Object Security](#)

Source and Target Host Security

Oracle on UNIX

Delphix support for Oracle on UNIX requires an OS account (**delphix_os**) on source database servers and on target servers that will host virtual databases or files. The Delphix Engine uses SSH to send commands to this user, which performs operations on the host. Some of these commands require elevated privileges.

 There is no actual requirement that the account be named **delphix_os** on both sources and targets. You can name the account anything you want; you can also use separate accounts on every source and target.

- **Restrict su - delphix_os to named Delphix Admins and System Administrators**

Other users of the system do not need access to the delphix_os user. Your Delphix Admins and System Administrators should retain su ability to facilitate troubleshooting.

- **Use SSH Key exchange to allow the Delphix Engine to communicate with targets**

Implement public/private key exchange instead of username/password. This allows you to keep the password of **delphix_os** completely secret.

- **Put delphix_os on password rotation**

Rotate the delphix_os password in accordance with your enterprise security policy for application software accounts. You should either:

- implement SSH Key exchange prior to placing delphix_os on password rotation, or
- script CLI commands to update the password inside the engine as part of the rotation process.

Delphix Professional Services can assist you in integrating Delphix with your enterprise password rotation system.

- **Restrict elevated privilege commands to the lowest level needed**

The Delphix Engine uses elevated privileges to provide core features as well as optional features. The Delphix docs describe in detail which privileges are absolutely necessary, as well as techniques for further restricting the commands that can be used.

The Delphix Engine ships with support for “sudo” as the privilege elevation system, but also allows for integration with third-party and custom centralized privilege management systems.

Windows

Delphix support for SQL Server requires two OS accounts for Windows:

- **delphix_src** – used on the source database server
- **delphix_trgt** – used on the servers which host Virtual Databases Both are required for the Validated Sync target

 There is no actual requirement that the account is named **delphix_src/delphix_trgt**. You can name the account anything you want; you can also use separate accounts on every source and target. Finally, you can create a single account for use everywhere, but this is not recommended since it violates the separation of duties.

Restrict privileged commands to the lowest level needed

The Delphix user or domain account should have exactly the privileges required in the Delphix documentation. Do not grant additional privileges.

Put delphix_src and delphix_trgt on password rotation

Change the user or domain account password at regular intervals or in accordance with security policies for application software accounts. Use CLI scripts to quickly modify the password across the Delphix ecosystem. Delphix Professional Services can assist you in scripting and integrating Delphix with your enterprise password rotation system.

Use minimum privileges on your SMB share

Consult <http://technet.microsoft.com/en-us/library/cc754178.aspx> to understand how shared folder privileges work. Use the minimum privileges.

Use Windows Authentication for SQL Server

SQL Server allows authentication via Windows or Mixed mode. Mixed mode allows authentication via Windows or SQL Server.

Windows authentication is more secure; it uses Kerberos security protocol, provides password policy enforcement with regard to complexity validation for strong passwords, provides support for account lockout, and supports password expiration. <http://msdn.microsoft.com/en-us/library/ms144284.aspx>.

Related Topics

- [Object Security](#)

System Configuration

There are a number of configuration options available in the System Administration area that help you to secure the Delphix Engine. You can configure these during installation through the setup wizard, or later by accessing the Delphix Setup screens.

Maintain System Time with NTP

Establish at least one, but preferably three, corporate NTP servers and sync your Delphix Engine to them. This ensures that audit and error messages display the correct time.

When configuring Delphix Engine on VMware, be sure to configure the NTP client on the ESX host to use the same servers that you enter here. On a vSphere client, the NTP client is set in the **Security Profile** section of the configuration process.

Enable Phone Home

Phone Home service will send critical information about the Delphix Engine to Delphix Support using HTTPS, on a periodic basis. The use of a Web Proxy Server is fully supported. Phone Home data allows Delphix Support to proactively detect Delphix Engines affected by critical vulnerabilities.

Register Your Delphix Engine

Registration is fast and easy, and you can do it with or without Internet connectivity from the Delphix Engine. Failing to register the Delphix Engine will impact its supportability and security in future versions.

Enable LDAP for Authentication

The LDAP protocol is used by enterprise authentication services. Enabling LDAP authentication allows your Delphix Engine to leverage the password control features of these products, such as expiration, lockout, and complexity.

Import your LDAP server certificate into your Delphix Engine, and enable SSL/TLS.

Enable SMTP and/or SNMP Monitoring

When the Delphix Engine encounters errors, it issues **alerts**. Configure SMTP and/or SNMP to forward **alerts** to your central monitoring system.

Related Topics

- [Security](#)

GUI Security

Overview

The sections in this article cover securing the Delphix GUI, which is similar to securing other web consoles. Some of these solutions include reducing the session timeout threshold, creating a signed certificate, and disabling HTTP access.

Reduce Inactive Session Timeout to 15 minutes

This means that a user will be booted from the session after 15 minutes of inactivity. This is done with a CLI command on a per-user basis by modifying the `sessionTimeout` property of the **User** object, as shown below. The default inactive timeout happens after 30 minutes.

```
myhost.delphix.com> cd user
myhost.delphix.com user> select delphix_admin
myhost.delphix.com user 'delphix_admin'> update
myhost.delphix.com user 'delphix_admin' update *> set sessionTimeout=15
myhost.delphix.com user 'delphix_admin' update *> commit
```

Use a URL from your Domain and Create a Signed Certificate

Do not use IP Addresses to access the Delphix Engine. Create a hostname and DNS entry, such as “delphix1.mycompany.com”. Delphix Support can assist in converting the engine from a self-signed certificate to a signed certificate that maps to your domain name. Please file a Support ticket to proceed.

Disable HTTP Access

Disabling HTTP or configuring HTTP to redirect connections to HTTPS is recommended to protect in-flight user credentials and connections with the engine. This can be done via the [command line](#) or through the [GUI](#).

Related Topics

- [Security](#)
- [Configuring HTTP Settings for the Delphix Engine](#)
- [CLI Cookbook: Changing HTTP and HTTPS Web Connections](#)

Masking Sensitive Data

Encryption does not protect data that is accessed through applications and database clients, the most likely attack vector. Masking sensitive data before it gets to non-production systems is a critical tool in the security arsenal.

Delphix provides an add-on masking product for simple cost-effective integration. It is also possible to integrate Delphix with any masking technology.

There are many topologies to consider, and the complete explanation of their pros and cons is outside the scope of this document. Delphix Professional Services can assist you in analyzing various masking solutions.

Related Topics

- [Security](#)

Audit Logs

Review Audit Logs Monthly

Conduct a monthly review of audit logs on your Delphix Engine. Pay particular attention to provisioning operations of unmasked databases, which creates new copies of your production data. See [Accessing Audit Logs](#) for instructions.

Forward Audit Logs to Central Server via Syslog

Forward audit logs to a central audit server using syslog techniques. Delphix Professional Services can assist you with scripts that facilitate this. See also [Setting Syslog Preferences](#) for configuration instructions.

Related Topics

- [Accessing Audit Logs](#)
- [Security](#)
- [Setting Syslog Preferences](#)

Support Security

At least every six months, regenerate the registration code and re-register the engine. See [Regenerating the Delphix Engine Registration Code](#) for instructions and the reasons why this is important for security.

Related Topics

- [Registration Management](#)
- [Security](#)

Delphix Operating System (DxOS)

Delphix Support accesses the DxOS for deep diagnostics and troubleshooting. Access to the Delphix Engine requires access to your network. Typically this is granted via shared troubleshooting sessions over Webex, with full transparency.

If desired, you can enable additional control so that access can only take place when you provide a token.

If you disable Support Access, do not have the token, and you are unable to login as a system administrator, it can become impossible for Delphix Support to login to repair your system. Example: the management stack crashes, the login system becomes unavailable, and you have disabled Support Access and do not have the token. For this reason, Delphix strongly recommends leaving Support Access enabled at all times. If you wish to disable access, generate a unique token once a month and place it in a secure location separate from the Delphix Engine.

Disable Support Access (Optional)

Support Access Control is managed as a system administrator in the Server Setup area.

- When set to DISABLED, it is impossible for Delphix Support to login to the DxOS.
- When set to ENABLED (the default), Delphix Support can login to the DxOS.
- When set to ENABLED and calendar time is set and a token generated, Delphix Support can only login with the token during that timeframe, which you provide. Generate the token once/month for an entire month and store it in a secure location separate from your Delphix Engine. When requested, provide the token through a secure means: in your support ticket, via email or SMS to a trusted entity, etc.

Related Topics

- [Security](#)

Password Policies

Getting Started

The password policy feature allows users to create their own custom password policies and enforce the password policy on non-LDAP Delphix Engine users.

Understanding Password Policies

A password policy is a named password policy that can be assigned to a user. It is a set of requirements that passwords must satisfy.

- `minLength` - A password must be longer than this length.
- `reuseDisallowLimit` - The user should not reuse old passwords. This tells the number of last used passwords disallowed to be reused as the new passwords.
- `uppercaseLetter` - A password must have at least one capital letter.
- `lowercaseLetter` - A password must have at least one lower case letter.
- `digit` - A password must have at least one digit.
- `symbol` - A password must have at least one symbol.
- `disallowUsernameAsPassword` - A password should not be the same as the user name.

Password policy requirements

When you set a password, it must differ from the most recent password and contain:

- at least 5 characters
- at least one uppercase letter
- at least one lowercase letter
- at least one numeric digit
- at least one symbol such as #, \$, !
- do not use username or reverse username

This policy applies to non-LDAP Delphix Engine users. This includes the default users, **delphix_admin** and **sysadmin**. The password policy does not apply to LDAP users.

Default password policy

By default, the Delphix Engine enforces the password policy named **NONE**, which enforces the least possible constraint.

Passwords must contain at least one character.

Changing the password policy

To change the current password policy from the default policy **NONE**, create a custom password policy and select it instead of **NONE**.

Who can change password policy for whom

- Domain administrators can change the current password policy for all domain users.
- System users can change the current password policy for all system users.
- Domain regular users (non-administrators) users can only view the password policy.

What operations can be done by administrators

- Create custom password policies
- Update custom password policies
- Delete custom password policies
- Change the current password policy to any of the available password policies
- View available password policies
- View current password policy requirements

Password policy parameters

When you create a password policy, you can set the following parameters:

- Unique name for the password policy
- Minimum length of the password
- Whether password must differ from the last password
- Whether password must not contain the username or reverse user name
- Whether password must contain at least one uppercase letter
- Whether password must contain at least one lowercase letter
- Whether password must contain at least one numeric digit
- Whether password must contain at least one symbol such as #, \$, !

Restrictions

- Restrictions for default password policy's modification (named **NONE**):
 - not allowed to delete the default password policy from available list of password policies.
 - not allowed to update any parameters of the default password policy.
- Cannot delete the password policy which is set as current password policy.

Related Links

- [CLI Cookbook: Password Policies](#)

Additional Topics

The Delphix Engine provides robust, enterprise-quality security controls. Performing the steps listed in this document will allow you to easily bring your Delphix Engines into compliance with your organization's security policies.

Perform a Yearly Audit

At least once annually, audit one or more Delphix Engines to ensure compliance with your security policies.

Port Scan

Delphix fully supports network security scans, using a tool of your choosing.

Security Testing

Many companies require security testing of applications in their environment using a Port Scanner or other Security Penetration Test tools. Delphix supports the use of these security tools with the application credentials available for the engine (e.g., `delphix_admin`). The Delphix Engine is a closed appliance, and OS credentials on the appliance are not provided for these tests.

Security Banner

You can configure a custom security banner that will be displayed to all the users prior to login.

1. Login to the CLI using the `sysadmin` **username** and **password**.

```
ssh sysadmin@yourdelphixengine
```

2. Set the configuration and commit the changes.

```
delphix > service security
delphix service security > update
delphix service security update * > set banner="Your Message Here"
delphix service security update * > commit
```

 The string set in the banner is in plain text only. You can also generate banners with multiple lines, however, this can only be done using the "securityConfig" API. Refer to the [Delphix KB article](#) for additional information.

Virtual Database Security

The Delphix Engine provides advanced storage capabilities and automation to allow rapid provisioning of virtual databases (VDBs), which use only a fraction of the physical storage used by full database copies. Nonetheless, a VDB is equivalent to a physical database **and must be properly secured like any other database**.

By far the most dangerous attack vectors in the Delphix ecosystem are the same ones that existed pre-Delphix: unauthorized access to your non-production systems containing sensitive production data. **You must perform all the same actions to harden virtual databases as you would to harden physical clones.**

For information on securing your virtual databases, consult vendor-specific material and security guides.

Related Topics

- [Security](#)

Datasets



Unknown macro: 'navbox'

Getting Started

In order to start using Delphix, you'll need to connect one or more data sources and appropriate hosts for an engine. These are key concepts you'll need to understand as an administrator using Delphix with any data source. An outline of the contents in this section:

This section covers the following topics:

- [Managing Environments and Hosts](#)
- [Managing Data Sources and Syncing Data](#)
- [Provisioning and Managing Virtual Databases](#)
- [Hook Scripts for Automation and Customization](#)
- [Managing Policies](#)
- [Accessing the Delphix Engine](#)

Managing Environments and Hosts

In Delphix, an environment is either a single instance host or cluster of hosts that run database software. For example, a Linux system running Postgres. This may either be an individual instance, or a cluster like Oracle RAC. Environments can either be a source (where data comes from), staging (where data are prepared/masked) or target (where data are delivered and used by developers and testers).

Each environment has its own properties and information depending on the type of environment it is. In this section, you can learn about how to add environments to Delphix as well as what attributes and configurations that environments have.

Environment Attributes and Configurations

This topic describes the attributes of an environment such as name, host address, ssh port, or toolkit path. For more advanced attributes for specific data platforms, please navigate to that data platform's documentation section.

Environment Attributes

Attribute	Description
Environment User(s)	The user(s) for the environment. These are the users who have permission to ssh into an environment (Unix-based environments) or access the environment through the Delphix Connector (Windows-based environments).
Host Address	The fully qualified domain name or IP address for the environment.
NFS Addresses (Unix-based hosts only)	A list of IP addresses to be used for NFS mount for this environment. If left empty, Delphix will automatically use the Host Address parameter above for NFS requests
Delphix Session Protocol (DSP) Configurations	If an environment already exists after enabling server/client DSP authentication, you will need to modify its attributes for host communication to continue working, you will need to set up the appropriate stores on the remote host.
SSH Port	Environment port used for SSH, usually port 22. This parameter is only used for Unix-based environments.
Toolkit Path	Location of the toolkit. For Windows hosts, this will be specified once installing the Delphix Windows Connector.

Java Path	Location of the Java Development Kit (JDK) used for the host. Only specified if the feature to provide your own JDK is enabled, otherwise, the defaults are used per our Java Support Matrix. (link) Delphix uses Eclipse Adoptium (formerly known as AdoptOpenJDK) by default. You can provide your own Java when adding an environment and in the environment configuration page,
Notes	Any other information you want to add about the environment

Environment Users

Not to be confused with Delphix administrator users and sysadmin users, environment users must be created for each environment to perform actions on each host. These users are configured for each data platform to interface with the database instance on the host.

1. Login to the Delphix Management application using Delphix Admin credentials.
2. Click **Manage**.
3. Select **Environments**.
4. Click on the existing environment name you want to modify and open the environment information screen.
5. In the **Details** tab, click the **Plus** icon located next to Environment users.
6. There are two ways that can be used for the Delphix Engine to login into the environment.
 - a. Enter the Username and Password for the OS user in that environment and click **Validate**.
 - b. If you want to use a [public key](#) for logging into your environment:
 - i. Select 'Username and Public Key' for the Login Type.
 - ii. Copy the public key that is displayed, and append it to the end of ~/.ssh/authorized_keys file of the new user being added. If this directory or file does not exist, you will need to create it.
 - iii. Run `chmod 600 authorized_keys` to enable only the file owner with read and write privileges.
 - iv. Run `chmod 755 ~` to make your home directory writable only by your user and no other user may write to it.
 - v. The public key needs to be added only once per user and per environment.
7. Click the **Add** icon to save the new user.
8. To change the primary user for this environment, select the environment. Then click the **'star'** icon next to **Environment Users**. Only the primary user will be used for environment discovery.
9. To delete a user, click the **Trash** icon next to their username.

Environment Operations

Environments managed by Delphix have two operations, refresh and delete.

1. **Environment Refresh**: will update the metadata associated with that environment and sends a new toolkit to the host.
2. **Environment Delete**: will remove all metadata associated with an environment on the engine.

Environment Refresh

After you make changes to an environment that you have already set up in the Delphix Management application, such as installing a new database home, creating a new database, or adding a new listener, you may need to refresh the environment to reflect these changes.

During environment discovery and environment refreshes, Delphix pushes a new copy of the toolkit to each host environment and will replace the old toolkit with the newer one. Included in the toolkit are:

- A Java Runtime Environment (JRE)
- Delphix jar files
- The HostChecker utility
- Scripts for managing the environment and/or VDBs
- Delphix Connector log files

When you refresh the environment it will push the toolkit back to the directory identified as the "Toolkit Path" for the given environment. Once this completes, you should be able to bring the dSource back online.

Delphix then executes some of these scripts to discover information about the objects in your environment (where the databases are installed, the database names, information required to connect to these databases, etc.). In some environments (Windows in particular), the scripts are customized to fit your environment.

Note

An environment refresh or discovery operation does not alter the source configuration of manually added databases. If you have added the databases manually, then Delphix does not update its source configuration upon discovering a change.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the Environments panel, click the name of the environment you want to refresh.
5. Select the **Refresh** icon.
6. In the Refresh confirmation dialog select **Refresh**.

Environment Delete

Deleting an environment in Delphix will remove all data and configurations related to that environment. This action only affects the environment metadata stored in the Delphix Engine. It will not affect your database installations or homes on the hosts or clusters that the environment is referencing.

Before you can delete an environment, you must first delete all dependencies such as dSources and virtual databases (VDBs)

1. Login to the **Delphix Management** application.
2. Select **Environments**.
3. In the Environments panel, select the environment you want to delete.
4. Click the Actions (...) menu, and select **Delete**.
5. In the Delete Environment confirmation dialog click **Delete**.

Using HostChecker to Validate Environments

HostChecker is a standalone program which validates that host machines are configured correctly before the Delphix Engine uses them as data sources and provision targets.

Please note that HostChecker does not communicate changes made to hosts back to the Delphix Engine. If you reconfigure a host, you must refresh the host in the Delphix Engine in order for it to detect your changes.

You can run the tests contained in the HostChecker individually, or all at once. You must run these tests on both the source and target hosts to verify their configurations. As the tests run, you will either see validation messages that the test has completed successfully, or error messages directing you to make changes to the host configuration.

HostChecker is distributed as a set of Java files and executables. You can find these files and executables in five distinct tarballs, each containing a different jdk corresponding to a particular platform (OS + processor). Together, these tarballs comprise the set of engines supported by Delphix.

When validating hosts during a new deployment, it is important to download the appropriate tarball for the host you are validating. Tarballs follow the naming convention "hostchecker_<OS>_<processor>.tar." For example, if you are validating a linux x86 host, you should download the tarball named hostchecker_linux_x86.tar.

HostChecker is also included in the Delphix Toolkit which is pushed to every environment managed by the Delphix Engine. It can be found in /<toolkit-path>/<Delphix_COMMON>/client/hostchecker.

Privilege Elevation Profiles

Privilege Elevation Profiles exist to provide the Delphix Engine with a mechanism for running privileged commands in a secure way to achieve the following:

- Mount and Unmount NFS filesystems
- Create and Remove directories in paths not owned by the Delphix OS user
- Examine the running process list
- Run commands as root

This is an advanced CLI topic and intended for advanced end-users and Delphix Professional Services consultants.

Support for Privilege Elevation Profiles

Writing and troubleshooting scripts, such as those required for Privilege Elevation Profiles, is out of scope and not covered by Delphix Support.

How do Privilege Elevation Profiles Work?

Privilege Elevation Profiles need to be tailor-made to work with non-standard environments that may use third-party or proprietary a privilege elevation mechanism other than sudo. Customers are strongly encouraged to work with Delphix Professional Services to formulate reliable profile scripts.

Privilege Elevation Profiles exist within a two-tier cascading hierarchy. This means there is one default profile for the entire Delphix Engine that should contain scripts for all the operations that require privilege elevation. Additional profiles may contain a subset of the scripts. When a non-default profile is used, the Delphix Engine uses that profile's scripts where they exist and reverts to the scripts in the default profile if no script for the operation exists. By default, the Delphix Engine ships with simple scripts that pass commands to the standard UNIX sudo command.

All Environments added to the Delphix Engine get added with the default Privilege Elevation Profile. The profile can be assigned on a per-host basis. Below shows how a host using a non-standard profile will use scripts in the cascading model.

In order to create a privilege elevation profile, you must create both a profile and a profileScript. Scripts exist for particular operations, which include:

- dlp_x_mount
- dlp_x_umount
- dlp_x_rmdir
- dlp_x_mkdir
- dlp_x_ps
- dlp_x_pexec

There are three parameters when creating a new profile:

1. name:
2. contents:

3. profile:

The Delphix Toolkit

The Delphix Toolkit is a group of files that allow communication between Delphix engines and remote hosts. Included in the Delphix Toolkit are:

- Database and Host Scripts
- Delphix Session Protocol (DSP) binaries
- Delphix HostChecker
- Java Development Kit (JDK)

The Delphix Toolkit is automatically sent to Environments at the specified Toolkit Path whenever adding or refreshing an environment.

Toolkit Size and Predicted Growth

Delphix Toolkit's size depends on two factors:

1. The number of Delphix client applications running on the environment
2. The number of VDBs on the environment (if any)

There are currently three Delphix client-side applications: SnapSyncClient, the Delphix Connector, and V2P. Each of the clients that run from the client-side toolkit generates their own logs - info, trace, debug, error. Each level of logging is restricted to a maximum of 10 log files and these log files are capped at 10 MB each.

Therefore the maximum space occupied by the toolkit directory on the Source server is its initial size of 400 MB + 1200 MB = 1.6 GB.

However, on the target server, the maximum toolkit size is initial size 400 MB + 800 MB + (Number of VDBs * 1MB)

Related Topics

- [Managing Oracle Environments and Hosts](#)
- [Managing SQL Server Environments and Hosts](#)
- [Managing SAP ASE Environments and Hosts](#)
- [Managing Db2 Environments and Hosts](#)
- [Managing PostgreSQL Environments and Hosts](#)
- [Managing HANA Environments and Hosts](#)

Managing Data Sources and Syncing Data

A dSource is a virtualized representation of a physical or logical source database, which cannot be accessed using database tools. In order to use dSource snapshots, you must create a virtual database (VDB), an independent, writable copy of a dSource snapshot.

Once an environment has been added, Delphix will automatically discover any data sources on the host. These appear as dSources on the Databases page of each environment. This section will describe how to manage these dSources and sync data from your sources. The following operations and concepts exist for dSources:

- Link/Sync
- Timeflows
- Snapshots
- SnapSync and LogSync
- Enable/Disable
- Detach/Reattach
- Delete
- Upgrade

Linking dSources

Once you have discovered your data source environments, you will see a list of data sources on that host. The first step to sync data is to link a dSource from the host.

Linking a dSource will ingest data from the source and create a dSource object on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

Once a dSource has been linked, you can begin to perform dSource operations on it, such as enable, detach, delete, and more.

The dSource Timeflow

Both dSources and VDBs have Timeflows, which are a collection of snapshots for a particular data set. Virtual databases can be provisioned from any snapshots in a Timeflow. In the UI, a Timeflow is represented vertically with a series of individual snapshots as rows.

The dSource Snapshot

Snapshots of dSources preserve specific points in time for use later during the virtual database provisioning workflow. Taking a snapshot will create a new snapshot entry in the dSource's Timeflow.

SnapSync and LogSync

This dSource Timeflow described above is maintained through the use of SnapSync and LogSync.

SnapSync

SnapSync will pull over the complete data set during the initial load using standard database protocols.

Subsequent SnapSync operations will pull only the incremental changes and store them with optimal storage efficiency. At the end of each SnapSync operation, a snapshot is created that serves as the base point for provisioning operations.

When provisioning a VDB, the closer the origin point is to a snapshot created via SnapSync, the sooner the provisioning operation will complete.

LogSync

In addition to SnapSync, LogSync will periodically connect to the host(s) running the source database via standard protocols and pull over any log files associated with the database. These log files are stored separately from the SnapSync data and are used to provision from points in between SnapSync snapshots.

When provisioning from a point between snapshots, the additional time it takes to provision is directly proportional to the time difference between the provision point and the last snapshot. The rate of change on the source database dictates the amount of data that must be replayed to bring a virtual database to the correct point in time.

Enabling and Disabling dSources

Some operations, such as dSource upgrade, are not available unless the dSource is disabled. Disabling a dSource turns off communication between the Delphix Engine and the source database, but it does not remove the configuration that enables communication and updates data. When a disabled dSource is later enabled, it will resume communication and incremental data updates from the source database according to the original policies and data management configurations that you set.

Disabling a dSource is also a prerequisite for several other operations, such as:

- database migration
- upgrading the dSource metadata after the upgrade of the associated data source
- restoring the source database from a backup

Disabling a dSource will stop further operations on the Delphix Engine related to the dSource.

1. Login to the **Delphix Management** application.
2. Select the dSource you want to disable.
3. In the upper right-hand corner, from the Actions menu (...) select **Disable**.
4. In the Disable dialog, select **Disable**.

When you are ready to **enable** the dSource again, from the Actions menu (...) select **Enable**, and the dSource will continue to function as it did previously.

Detaching and Reattaching dSources

Each dSource contains metadata that associates it with the source database, as well as the data it has ingested from the source database in the form of snapshots up to that point.

You may want to detach a dSource from a data source and reattach it to another source. It is possible to detach, or unlink, a dSource from its source database. This breaks the association with the source database without affecting the data within the Delphix Engine.

Detached dSources and their source databases have these properties:

- A detached dSources can still be used to provision a virtual database (VDB).
- You can re-link the source database as a different dSource.
- Any child VDBs that were provisioned from this dSource will only be able to be refreshed from the most recent snapshot available on the dSource.
- If you need a VDB from a newer snapshot, you will need to provision a new VDB. Once you have provisioned the new VDB you can delete the old VDBs provisioned from this dSource. You can delete the old dSource when it is no longer needed.

Detaching a dSource

1. Login to the **Delphix Management** application.

2. Select the dSource you want to detach.
3. From the Actions menu (...), select **Unlink dSource**.
A warning message will appear.
4. Click **Unlink** to confirm. The status of the dSource will show as Detached.

When you are ready to **attach** the dSource again, from the Actions menu (...) select **Link dSource**, and the dSource.

Deleting dSources

Deleting a dSource will delete the dSource metadata for a particular source database, along with all associated snapshots, logs, and policies stored in Delphix. If a dSource is deleted, it will not affect your source database.

Prerequisites

You cannot delete a dSource that has dependent virtual databases (VDBs). Before deleting a dSource, ensure that you have deleted all dependent VDBs as described in *Deleting a VDB* on [Provisioning and Managing Virtual Databases](#). Otherwise, the delete option will be grayed out.

Note

For dSources, you cannot delete a dSource that has dependent virtual databases (VDBs) or virtual Pluggable databases (vPDBs). Before deleting a dSource, ensure that you have deleted all dependent VDBs or vPDBs. Otherwise, the delete option will be grayed out. When you delete all the dSources of PDBs, the associated CDBs are automatically deleted.

Procedure

1. Log in to the **Delphix Management** application.
2. Select the dSource you want to delete.
3. From the Actions menu (...), select **Delete**.
4. Click **Delete** to confirm.

Upgrading dSources

Occasionally, you may upgrade the source database that a dSource is associated with. In this case, you will need to update the dSource's metadata to associate a dSource with a more recent database version, thus performing an 'upgrade' of the dSource's known version.

While this feature does not upgrade the actual database, it is important that Delphix associates this dSource with the correct version of the source database to function properly.

This process of upgrading a dSource will modify the database installation path to use the path of the newer database installation.

The process for the upgrade of dSources will vary depending on the database. Refer to the following pages for database-specific procedures:

- [Upgrading dSources after an Oracle Upgrade](#)
- [Upgrading a dSource after a SQL Server Upgrade](#)
- [Upgrading a SAP ASE dSource](#)

Prerequisites

- The data source instance has been upgraded to a more recent installation version.
- The dSource must be disabled.

Related Topics

- [Getting Started](#)
- [Linking Data Sources and Syncing Data with Oracle](#)
- [Linking Data Sources and Syncing Data with SQL Server](#)
- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Linking Data Sources and Syncing Data with Db2](#)
- [Linking Data Sources and Syncing Data with PostgreSQL](#)
- [Linking Data Sources and Syncing Data with HANA](#)

Provisioning and Managing Virtual Databases

After you have added a data source from an environment and created snapshots of a dSource, you can begin to provision virtual databases (VDBs), one of the key parts of using Delphix.

As discussed before, a dSource is a virtualized representation of a physical or logical source database, which cannot be accessed or manipulated using database tools. In order to use dSource snapshots, you must create a VDB, an independent, writable copy of a dSource snapshot.

You can create VDBs from other VDBs. Once a VDB has been provisioned to a target environment, you can also create a snapshot policy for that VDB, to capture changes within it as if it were any other logical or physical database.

The following operations and concepts exist for VDBs:

- [Provision](#)
- [Start/Stop](#)
- [Refresh](#)
- [Rewind \(Rollback\)](#)
- [Enable/Disable](#)
- [Delete](#)
- [Force Delete](#)
- [Migrate](#)
- [Upgrade](#)
- [Virtual to Physical \(V2P\)](#)
- VDB Configuration Templates

Provisioning Virtual Databases

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment, as described in the overview for [Managing Environments and Hosts](#). Since provisioning VDBs varies significantly across data platforms, details on this process can be found within the documentation for a specific data platform.

At their core, virtual databases are fully functional copies of data sources that can be created and managed at a fraction of the storage and time that are typically required. To end-users, VDBs act and perform just like a standard database.

The VDB Timeflow

Both dSources and VDBs have timeflows, which are a collection of snapshots for a particular data set. Virtual databases can be provisioned from any snapshots in a Timeflow. The Virtual database's filesystem will be dependent upon the snapshot used to Provision, and as such, that snapshot can not be removed until the dependency is removed. In the product UI, a Timeflow is represented vertically with a series of individual snapshots as rows.

The VDB Snapshot

Snapshots of dSources preserve specific points in time for use later during the virtual database provisioning workflow. Taking a snapshot will create a new snapshot entry in the VDBs Timeflow.

Refreshing Virtual Databases

Refreshing a VDB will re-provision it from the dSource. As with the normal provisioning process, you can choose to refresh the VDB from a snapshot or a specific point in time. However, you should be aware that refreshing a VDB will delete any changes that have been made to it over time. When you refresh a VDB, it will revert to its past state in the specified snapshot or point in time.

When performing a VDB refresh, there is an option to choose between either the faster or more accurate point in time for that database.

i Although the VDB no longer contains the previous contents, the previous Snapshots and TimeFlow still remain in Delphix and are accessible through the Command Line Interface (CLI).

Prerequisites

To refresh a VDB, you must have the following permissions:

- **PROVISIONER** permissions on the dSource associated with the VDB
- **PROVISIONER** permissions on the group that contains the VDB
- **Owner** permissions on the VDB itself
- **Data** is a role that allows DB_ROLLBACK, DB_REFRESH, READ_ACTION, DB_SYNC, JOB_CANCEL.
- **Read** is a role that allows the user to inspect objects via the READ_ACTION permission.

A user with admin credentials can perform a VDB Refresh on any VDB in the system.

Procedure

Perform the following steps to refresh a VDB.

1. Log in to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Select the VDB that you want to refresh.
4. Click the **Refresh VDB** button. The Refresh VDB screen appears.



5. Select one of the refresh options and click **Next**.
 - a. **Faster** - This option will utilize the most recent snapshot in the timeflow.
 - b. **More Accurate** - This option allows you to select a snapshot, a point in time, or a specific log ID.
6. Click **Submit** to confirm.
7. Click the **Actions** button to watch the progress of the refresh job.
8. If you want to know when the VDB was last refreshed/provisioned, check the Time Point on the **Status** page.

Rewinding Virtual Databases

Rewinding a VDB rolls it back to a previous point in its Timeflow and re-provisions the VDB. The VDB will no longer contain changes after the rewind point.

i **Rewind VDB**

Although the VDB no longer contains changes after the rewind point, the rolled-over Snapshots and Timeflow still remain in Delphix and are accessible through the Command Line Interface (CLI). See the topic [CLI Cookbook: Rolling Forward a VDB](#) for instructions on how to use these snapshots to refresh a VDB to one of its later states after it has been rewound.

i Delphix clones a new Timeflow from the closest Snapshot older than or equal to the rewind point. This creates a dependency between the new Timeflow and the parent Snapshot and Timeflow. The parent Snapshot and Timeflow cannot be deleted because of this dependency. The VDB must first be refreshed before the parent Snapshot and Timeflow can be removed.

Prerequisites

To rewind a VDB, you must have the following permissions:

- **Auditor** permissions on the dSource associated with the VDB
- **Owner** permissions on the VDB itself

You do NOT need **owner** permissions for the group that contains the VDB. A user with admin credentials can perform a VDB Rewind on any VDB in the system.

Procedure

1. Log in to the **Delphix Management** application.
2. Select the VDB you want to rewind.
3. Select the **Timeflow** tab.
4. Select the rewind point as a snapshot or a point in time.
5. Click **Rewind**.
6. If you want to use login credentials on the target environment other than those associated with the environment user, click **Provide Privileged Credentials**.
7. Click **Rewind** to confirm.

i **Using Timeflow Bookmarks**

You can use TimeFlow bookmarks as the rewind point when using the CLI. Bookmarks can be useful to:

- Mark where to rewind to - before starting a batch job on a VDB for example.
- Provide a semantic point to revert back to in case the chosen rewind point turns out to be incorrect.

For a CLI example using a TimeFlow bookmark, see [CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark](#).


Undoing a Rewind or Refresh

An accidental refresh of one or more virtual databases (VDBs) from a source database has removed important data. To correct this situation users can undo VDB refresh or rewind actions.


Procedure

To undo a refresh or rewind via the Delphix Management application:


1. Login to the **Delphix Management** application.
2. Under **Datasets**, select the **VDB** you want to undo your rewind/refresh for.
3. From the **Actions** menu (...) select **Undo Refresh or Rewind**.



4. Confirm that you want to undo the last refresh or rewind for the selected VDB, by selecting **Undo**.




From the Action sidebar menu, you will be able to view the undo action.



Starting and Stopping Virtual Databases

Stopping your virtual databases is necessary if you are performing any actions on the database that require it to be offline. In addition to properly stopping the database, the Delphix stop and start actions also unmount and mount the filesystems that are provided by the Delphix Engine as well as performing any Delphix-specific actions during startup or shutdown.

1. Login to the **Delphix Management** application.
2. Select the VDB you want to manage.
3. To stop an active VDB, click the **stop** icon in the top right-hand corner.
To start an inactive VDB, click the **start** icon in the top right-hand corner.
4. If stopping or starting the VDB requires particular credentials for the target environment other than those of the default environment user:
 1. Check Provide Privileged Credentials.
 2. Enter the username and password.
 3. Click Validate Credentials.
 - Click Yes to stop or start the VDB.
 - You can view the status of the stopping or starting of the VDB in several ways:
 - View the **Active Jobs** on the right of the screen or monitor the progress bar from the VDB Status tab.



Enabling and Disabling Virtual Databases

Disabling a VDB is a prerequisite for procedures such as VDB migration or upgrade. Disabling a VDB removes all traces of it, including any configuration files, from the target environment to which it was provisioned. However, the VDB itself, as well as the metadata, will still exist on the engine.

When you later enable the VDB again, these configuration files are restored on the target environment.

1. Login to the **Delphix Management** application.
2. Select the VDB you want to disable.
3. From the Actions menu select **Disable**.
4. Click **Disable** to acknowledge the warning.

When you are ready to enable the VDB again, from the Actions menu select **Enable**, and the VDB will continue to function as it did previously.

Deleting Virtual Databases

Deleting a VDB will remove it and its Timeflow and snapshots from the Delphix Engine entirely. This operation is non-reversible.

Prerequisites

You cannot delete a VDB that has dependent VDBs. Before deleting a VDBs, make sure that you have deleted all dependent VDBs. Otherwise, the delete option will be grayed out.

 Deleting a VDB is an unrecoverable operation. Proceed only if you want to permanently destroy the unique data that was created in the VDB.

Procedure

1. Login to the **Delphix Management** application.
2. Select the VDB you want to delete.
3. From the **Actions** menu select Delete.
4. Click **Delete** to confirm.

Force Delete


Deleting a vFiles may fail if it cannot be unmounted successfully from all target environments. You can use the **Force Delete** option to ignore all failures during unmount.

Using Force Delete

Deleting unused or outdated objects should be a regular part of Delphix Engine administration. This is especially important to prevent low space errors, which can cause the Delphix Engine to stop. **The Delphix Engine holds a maximum of 400 objects.**

1. Log into the **Delphix Management** application.
2. Select **Resources > Storage Capacity**.
3. Next to the object you want to delete select the **Trashcan**.

- In the Delete dialog select **Force Delete**. Oracle users will have the option to provide additional credentials.



Delete Dataset Child VDB [X]

Are you sure you want to delete dataset "Child VDB"?

Force Delete


Provide privileged credentials

Cancel Delete

- Click Delete.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot. These items are displayed with a lock icon next to the name.



Migrating Virtual Databases

There may be situations in which you want to migrate a virtual database to a new target environment, for example when upgrading the host on which the VDB resides, or as part of a general data center migration. Using the VDB migrate functionality allows you to change the VDB's location from one target environment to another.

This is easily accomplished by first disabling the database, then using the Migrate VDB feature to select a new target environment.

Prerequisites

- You should have already set up a new target environment that is compatible with the VDB that you want to migrate.
- Ensure that the VDB has first been disabled

Procedure

1. Login to the **Delphix Management** application.
2. Select the VDB you want to migrate.
3. From the Actions menu select **Migrate**.
4. Select the new target environment for the VDB, the user for that environment, and the database installation where the VDB will reside.
5. Select **Migrate** to confirm your selections.
6. From the **Actions** menu select **Enable**.
7. Click **Enable** to confirm.

Within a few minutes, your VDB will re-start in the new environment, and you can continue to work with it as you would any other VDB.

Upgrading Virtual Databases

Similar to upgrading a data source, upgrading a VDB changes the installation path that is associated with that data platform. While this feature does not upgrade the actual VDB, it is important that Delphix associates this VDB with the correct database version to function properly.

By performing this upgrade, you can update the version of the database installation used for running the VDB.

Procedure

1. Login to the **Delphix Management** application.
2. Select the VDB you want to upgrade.
3. From the **Actions** menu select **Upgrade**.

Virtual to Physical (V2P)

Virtual to Physical, or V2P, is the process of exporting a virtual database to a physical one.

After you have created a dSource or a VDB, you can export its contents and log files to a physical database. This process creates a set of directories on the target environment and populates them with the database data, log files, and scripts that are used to recover the physical database.

You can automatically start the physical database recovery process as part of V2P, or you can use the scripts for manual recovery. When the export process completes, the target environment will contain a copy of the database in its unvirtualized size, so before you begin the process, make sure the target directories you specify during V2P have enough capacity to hold the database. For more information, visit the following page: [Virtual to Physical](#).

Procedure

1. Login to the **Delphix Management** application.
2. Select the dSource or VDB you want to export.
3. Select the snapshot of the dSource or VDB state you want to export.
4. From the **Actions** menu select **Virtual to Physical**.
5. Select the target environment, and fill the other input fields depending on the data platform of your object.

Provisioning with Replication

This topic describes how to provision from a replicated dSource or virtual database (VDB).

The process for provisioning from replicated objects is the same as the typical VDB provisioning process, except that first you need to select the replica namespace containing the replicated object.

Prerequisites

- You must have a dSource or VDB that has been replicated from one Delphix Engine to another, as described in [Replication Overview](#).
- The Delphix Engine containing the replicated dSource or VDB must have a compatible target environment that it can use to provision a VDB from the replicated dSource or VDB.
- On the Delphix Engine containing the replicated dSource or VDB, login to the Delphix Management application.

Procedure

1. From the list of replica namespaces, select the replica namespace that contains the dSource or VDB from which you want to provision.
2. The provisioning process is now identical to the process for provisioning standard objects.

Related Topics

- [Provisioning and Managing Virtual Databases with Oracle](#)
- [Provisioning and Managing Virtual Databases with SQL Server](#)
- [Provisioning and Managing VDBs from SAP ASE](#)
- [Provisioning and Managing VDBs from Db2 dSources](#)
- [Provisioning and Managing VDBs from PostgreSQL dSources](#)
- [Provisioning and Managing VDBs from HANA dSources](#)

Virtual to Physical

This section covers the following topics:

- [Customizing Target Directory Structure for Database Export](#)
- [Manually Recovering a Database After V2P](#)
- [V2P with Unstructured Files](#)

Customizing Target Directory Structure for Database Export

This topic describes how to customize the target directory layout for database export.

In the V2P export process, it may be necessary to customize the target directory structure which the files will be exported to. The following is the default directory structure:

- Data files: <target directory>/data
- Archive files: <target directory>/archive
- Temp files: <target directory>/temp
- External files: <target directory>/external
- Script files: <target directory>/script

Note: The example on this page uses / for file separators which is relevant for Unix and Linux environments. If the target environment is Windows, the file separator will be \ .

The following procedure describes how to customize the directory layout.

Procedure

1. During the virtual to physical export process, click **Advanced** in the V2P Wizard to see the target directory options.
2. You can customize any of the following:
 - **Data Directory**
 - **Archive Directory**
 - **Temp Directory**
 - **External Directory**
 - **Script Directory**
3. Each directory will then be concatenated to the **Target Directory** separated by the appropriate separator.

i Any one of **Target Directory, Data Directory, Archive Directory, Temp Directory, External Directory, Script Directory** can be blank. However, the combination of the fields must form an absolute path.

- Data files: <target directory>/<data directory>
- Archive files: <target directory>/<archive directory>
- Temp files: <target directory>/<archive directory>
- External files: <target directory>/<external directory>
- Script files: <target directory>/<script directory>

Examples

Target directory is not empty

This means all target directories have a common root.

Input	Final Directory Layout
<ul style="list-style-type: none"> • Target Directory: /mytarget • Data Directory: /mydata • Archive Directory: /myarchive • Temp Directory: /mytemp • External Directory: /myexternal • Script Directory: /myscript 	<ul style="list-style-type: none"> • Data files: /mytarget/mydata • Archive files: /mytarget/myarchive • Temp files: /mytarget/mytemp • External files: /mytarget/myexternal • Script files: /mytarget/myscript

Target directory is empty

All target directories may not have a common root. Note that external files and temp files share the same common root.

Input	Final Directory Layout
<ul style="list-style-type: none"> • Target Directory: • Data Directory: /mydata • Archive Directory: /myarchive • Temp Directory: /mytarget/temp • External Directory: /mytarget/external • Script Directory: /myscript 	<ul style="list-style-type: none"> • Data files: /mydata • Archive files: /myarchive • Temp files: /mytarget/temp • External files: /mytarget/external • Script files: /myscript

Target directory is empty and data directory /

Combined with [Customizing VDB File Mappings](#), exporting data files to separate file systems is possible. In this example, `a.dbf` and `b.dbf` can be exported to `/filesystem1` and `/filesystem2` respectively. For Oracle VDBs, it is necessary to specify the parameter `useAbsolutePathForDataFiles` in order to export data files to separate file systems. See [CLI Cookbook: V2P \(Virtual to Physical\) of a Single Instance Oracle Database with Datafiles on Separate File Systems](#) for details.

Input	Final Directory Layout
<ul style="list-style-type: none"> • Target Directory: • Data Directory: • Archive Directory: /myarchive • Temp Directory: /mytarget/temp • External Directory: /mytarget/myexternal • Script Directory: /myscript • File mappings: <ul style="list-style-type: none"> • a.dbf : /filesystem1/a.dbf • b.dbf : /filesystem2/b.dbf 	<ul style="list-style-type: none"> • Data files: <ul style="list-style-type: none"> • /filesystem1/a.dbf • /filesystem2/b.dbf • Archive files: /myarchive • Temp files: /mytarget/mytemp • External files: /mytarget/myexternal • Script files: /myscript

Target directory is empty and one of the subdirectories is empty would result in an error.

Input	Final Directories
<ul style="list-style-type: none"> • Target Directory: • Data Directory: /mydata • Archive Directory: • Temp Directory: /mytarget/temp • External Directory: /mytarget/myexternal • Script Directory: /myscript 	<ul style="list-style-type: none"> • INVALID

 For parallel execution of V2P operations from the same target, make sure that the scripts directory name is unique.

Related Topics

- [Virtual to Physical](#)

Manually Recovering a Database After V2P

This topic describes how to manually recover a database after the V2P process.

⚠ For Oracle database, while running V2P you can set `recoverDatabase=false` and `openDatabase=false`. If you set `recoverDatabase=false`, you need to run `recover-vdb.sh`, followed by `open-vdb.sh`. But if you perform V2P with `recoverDatabase=true` and `openDatabase=false`, you will need to run `open-vdb.sh`.

Procedure

1. In the V2P target environment, navigate to the scripts directory for your exported database instance. You can find the scripts in a sub-directory named for that specific database instance. For Oracle databases, the path is `<target_directory>/<db_unique_name>/script/<instance name>`. For SQL Server databases, the path is `<target_directory>\<db_name>\scripts`.
2. For **Oracle** databases, locate the scripts `recover-vdb.sh` and `open-vdb.sh`. Run them in that order. For **SQL Server** databases, locate the script `Provision.ps1` and run it.
3. For **SQL Server** databases, when the script completes, **Refresh** the target environment for it to discover the recovered database.
For **Oracle** databases, add the recovered database to `/etc/oratab` and **Refresh** the target environment for it to discover the recovered database.

Related Topics

- [Virtual to Physical](#)

V2P with Unstructured Files

This topic describes the procedure for performing virtual to physical (V2P) operations with unstructured files.

V2P Not Supported for Unstructured Files on Windows

V2P is not supported for unstructured files on Windows environments. Similar results to V2P may be achieved by provisioning a vFiles and copying data out of the vFiles to the local machine.

Procedure

1. Log into the **Delphix Management** application.
2. Select the dataset you want to export.
3. Select the snapshot you want to export.
4. From the **Actions (...)** menu select **Virtual to Physical**.
5. Select the target environment.
6. Enter the **Mount Path** for the export.
The directory you enter here must exist in the target environment, and you must have permission to write to it.
7. Click **Next**.
8. Review the **Target Environment** configuration information, and then click **Submit**.

Related Topics

- [Virtual to Physical](#)

Hook Scripts for Automation and Customization

Hook operations allow you to execute custom operations at select points during linking sources and managing virtual datasets. For example, you may want to prevent your monitoring systems from triggering during VDB startup and shutdown. In this case, you would leverage pre- and post-hooks to run required scripts for VDB start/stop operations.

Unix-based operating systems will execute hook scripts in bash, while Windows operating systems will execute using PowerShell.

Each Delphix object will have its own set of hooks you can create. To easily manage hooks across multiple datasets, you can create hook templates to apply the same hook to many objects.

Setting Hook Operations

You can construct hook operation lists through the Delphix Management application or the command-line interface (CLI). You can either define the hooks as part of the provisioning process or edit them on virtual datasets that already exist.

Hook operations can be created in the Hooks tab of the Add dSource or Add VDB wizards.

To edit hook operations on a dataset, navigate to the datasets page and find the object's configuration. You will see a 'Hooks' tab where you can manage and modify hook operations.

Hook Operation Templates

You can use templates to store commonly used operations, which allows you to avoid repeated work when an operation is applicable to more than a single virtual dataset. You manage templates through the Delphix Management application.

Hook Operations List

dSource Hooks

Hook	Description
Pre Sync	Pre Sync hook operations are executed before a Snapsync of a dSource. These operations can quiesce data to be captured during the Snapsync, or stop application processes that may interfere with the Snapsync operation.
Post Sync	Post Sync hook operations are executed after a Snapsync operation of a dSource. These operations can undo any changes made by the pre-sync hooks. They will run regardless of the success of pre-sync hook operations or the Snapsync itself.

Virtual Database Hooks

Hook	Description
------	-------------

Configure Clone	<p>Operations performed after initial provision or after a refresh.</p> <p>This hook will run after the virtual dataset has been started.</p> <p>During a refresh, this hook will run before the Post Refresh hook.</p>
Pre Refresh	<p>Operations performed before a refresh.</p> <p>This hook will run before the virtual dataset has been stopped.</p> <p>These operations can cache data from the virtual dataset to be restored after the refresh completes.</p>
Post Refresh	<p>Operations performed after a refresh.</p> <p>This hook will run after the virtual dataset has been started and after the Configure Clone hook.</p> <p>This hook will not run if the refresh or Pre Refresh hook operations fail.</p> <p>These operations can restore cached data after the refresh completes.</p>
Pre Rollback	<p>Operations performed before a rollback, also known as rewind.</p> <p>This hook will run before the virtual dataset has been stopped.</p> <p>These operations can cache data from the virtual dataset to be restored after the rewind completes.</p>
Post Rollback	<p>Operations performed after a rollback, also known as rewind. This hook will not run if the rewind or Pre Rewind hook operations fail.</p> <p>This hook will run after the virtual dataset has been started.</p> <p>This hook will not run if the rewind or Pre Rewind hook operations fail.</p> <p>These operations can restore cached data after the rewind completes.</p>
Pre Snapshot	<p>Operations performed before a snapshot.</p> <p>These operations can quiesce data to be captured during the snapshot, or stop processes that may interfere with the snapshot.</p>
Post Snapshot	<p>Operations performed after a snapshot.</p> <p>This hook will run regardless of the success of the snapshot or Pre Snapshot hook operations.</p> <p>These operations can undo any changes made by the Pre Snapshot hook.</p>
Pre Start	<p>Operations performed before the startup of a VDB or vFile.</p> <p>These operations can be used to initialize configuration files or stop processes that might interfere with the virtual dataset.</p>

Post Start	Operations performed after the startup of a VDB or vFile. These operations can be used to clean up any temporary files, or restart processes that may have been stopped by a Pre-Start hook, or log notifications.
Pre Stop	Operations performed before the shutdown of a VDB or vFile. These operations can quiesce data or processes prior to the virtual dataset shutdown.
Post Stop	Operations performed after the shutdown of a VDB or vFile. These operations can be used to log notifications, clean up any temporary files, or stop/restart related processes.

Staging Server Hooks

Staging Server hooks are currently applicable to SAP ASE and staged plugins (Db2, HANA, and PostgreSQL) environments, and execute on the configured staging server.

Hook	Description
Pre Validated Sync	Operations performed on a staging server before validated sync. This hook will run before the validated sync operation whenever the validated sync run is triggered.
Post Validated Sync	Operations performed on staging server after validated sync. This hook will run after the validated sync operation whenever the validated sync run is triggered.

Related Topics

- [Setting Hook Operations](#)
- [Hook Operation Templates](#)

Setting Hook Operations

You can construct hook operation lists through the Delphix Management application or the command-line interface (CLI). You can either define the operation lists as part of the provisioning process or edit them on virtual datasets that already exist.

Setting Hook Operations through the Delphix Management Application

 The provisioning wizard still imports hook operations from templates.

Hook operations can be provisioned in the **Hooks** tab of the Add dSource or Add VDB wizards.

1. Select the **type of operation** and enter a name, operation type, and script.
2. To remove an operation from the list, click the **Trash** icon on the operation.
3. When you have set all hook operations, click **Next** to continue with the provisioning process.

To edit hook operations on a virtual dataset:

 From the **Datasets** panel, you can create hook operations from a template.

1. In the **Datasets** panel, click the virtual dataset.
2. Click the **Configuration** tab.
3. Within the **Configuration** tab, click the **Hooks** tab.
4. Select the **hook** to edit.
5. The current operations at this hook will be displayed. To edit this list of operations, click the **Pencil** icon in the top right-hand corner of the tab.
6. Click the **Plus** icon to add a new operation.
7. Select the **type of operation** or click  to load a hook operation template.
8. Click the **text area** and edit the contents of the operation.
9. To remove an operation from the list, click the **Trash** icon on the operation.
10. When you have set all hook operations, click the **checkmark** to save the changes.

 The provisioning wizard still imports hook operations from templates.

Setting Hook Operations through the CLI

To specify hook operations during linking, edit the relevant hook's array of operations defined on the `LinkingParameters > Source > Operations` object.

To specify hook operations during provisioning, edit the relevant hook's array of operations defined on the `ProvisionParameters > Source > Operations` object.

To edit hook operations on a dSource that already exists, edit the relevant hook's array of operations defined on the `Source > Operations` object.

To edit hook operations on a virtual dataset that already exists, edit the relevant hook's array of operations defined on the `Source > Operations` object.

For more information about these CLI objects, see the following documentation in the **Help** menu of the **Delphix Management** application:

- `LinkedSourceOperations`

- `VirtualSourceOperations`
- `RunCommandOnSourceOperation`
- `RunExpectOnSourceOperation` API

Example of Editing Hook Operations through the CLI

1. Navigate to the relevant source's `VirtualSourceOperations` object.
2. Select a **hook** to edit.

```
delphix> source
delphix source> select "pomme"
delphix source "pomme"> update
delphix source "pomme" update *> edit operations
delphix source "pomme" update operations *> edit postRefresh
```

3. Add an operation at index 0.

```
delphix source "pomme" update operations postRefresh *> add
delphix source "pomme" update operations postRefresh 0 *> set
type=RunCommandOnSourceOperation
delphix source "pomme" update operations postRefresh 0 *> set command="echo
Refresh completed."
delphix source "pomme" update operations postRefresh 0 *> ls
Properties
  type: RunCommandOnSourceOperation (*)
  command: echo Refresh completed. (*)
delphix source "pomme" update operations postRefresh 0 *> commit
```

4. Add another operation at index 1 and then delete it.

```
set command="echo Refresh completed."delphix source "pomme" update operations
postRefresh *> add
delphix source "pomme" update operations postRefresh 1 *> set
type=RunCommandOnSourceOperation
delphix source "pomme" update operations postRefresh 1 *> set command="echo
Refresh completed."
delphix source "pomme" update operations postRefresh 1 *> back
delphix source "pomme" update operations postRefresh *> unset 1
delphix source "pomme" update operations postRefresh *> commit
```

Related Topics

- [Hook Scripts for Automation and Customization](#)
- [Passing Credentials Securely to Hook Operations](#)
- [Hook Operation Templates](#)

Passing Credentials Securely to Hook Operations

Sometimes, commands in hook operations need credentials to perform tasks such as making API calls or managing local services. To avoid hard-coding credentials in the hook code, you can pass credentials securely to hooks via environment variables using a new list property of each hook operation. The elements in this list define a set of environment variables and the credentials to pass to the hook in those variables.

In the browser, the list can be viewed and updated in the **Credential Environment Variables** section of a hook operation. (Hook operations can be accessed in the dataset and hook template creation dialogs and view/edit pages.) For example, to view or edit hook credentials for a dataset:

1. Login to the **Delphix Management** application.
2. Under **Datasets**, select a **dSource** or **VDB**.
3. Select the **Configuration** tab and then the **Hooks** tab.
4. Select a hook operation.
5. To make any changes, such as adding, modifying or deleting credentials, click on the edit icon to the right of **Operation Type**.

In the API and CLI, the credentials list property of each hook operation is an array called `credentialsEnvVarsList`.

Hook credentials can be configured directly as a password or key (entered by the user) or indirectly by selecting a password vault and a location in that vault. Credentials entered directly are managed securely by the engine in the same way as other passwords and keys, which are encrypted in disk and neither logged nor exposed by the API, whereas vault credentials are managed securely and externally to the engine by a specialized product such as CyberArk or HashiCorp Vault that has already been [configured in the engine](#). Credentials stored in a vault are always retrieved by the engine just before executing the hook, so they will always contain the latest values. Any number of credentials of any type can be defined for an operation in the `credentialsEnvVarsList` property.

Four types of credentials can be defined for hooks:

- `CyberArkVaultCredential`
- `HashiCorpVaultCredential`
- `PasswordCredential`
- `KeyPairCredential` (supported only via the API and CLI at the moment)

For more information on each of these types, visit your engine's API page for that type at <https://<engine address>/api/#<type name>> (you must log in first).

To illustrate, consider a hook that needs to authenticate to two different APIs, one using an API key entered directly by the user and another using a username and password managed by a vault. The user can configure the credential variables as follows:

```

delphix source "pomme" update operations postRefresh *> edit
0.credentialsEnvVarsList
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
*> add
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
0 *> set baseVarName=API1
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
0 *> edit credentials
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
0 credentials *> set type=PasswordCredential
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
0 credentials *> set password="API-KEY-02a0b73f"
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
0 credentials *> back
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
*> add
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 *> set baseVarName=API2
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 *> edit credentials
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 credentials *> set type=CyberArkVaultCredential
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 credentials *> set
queryString="Safe=Apis;Folder=Root\Service2;Object=ApiClient"
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 credentials *> set vault=MyCyberArk1
delphix source "pomme" update operations.postRefresh 0 credentialsEnvVarsList
1 credentials *> commit

```

When this hook executes, it will receive two sets of environment variables containing the corresponding credentials: `API1_PASSWORD` with value `API-KEY-02a0b73f` and variables `API2_USER` and `API2_PASSWORD` containing the values stored in the vault `MyCyberArk1` at the location of the provided query string. The `API2` credentials are retrieved from the vault just before the hook executes, so they will always contain the latest values.

The complete list of environment variables that may be available to a hook is:

- `<base name>_USER` : user name. Present only if the credentials are of a vault credential type and the entry in the vault contains a user name.
- `<base name>_PASSWORD` : a password. Present for `PasswordCredential` . Also present in case of a vault credential type and the vault entry contains a password.
- `<base name>_PRIVKEY` : a private key. Present for `KeyPairCredential` . Also present in case of a vault credential type and the vault entry contains a private key.
- `<base name>_PUBKEY` : a public key. Present for `KeyPairCredential` . Also present in case of a vault credential type and the vault entry contains a public key.

Related Topics

- [Hook Scripts for Automation and Customization](#)
- [Password Vault Support](#)

Hook Operation Templates

You can use templates to store commonly used operations, which allows you to avoid repeated work when an operation is applicable to more than a single virtual dataset. You manage templates through the Delphix Management application.

Hook Operations Templates Not Available via CLI

Hook operation templates cannot be fully utilized from the CLI. Manage and use hook operations through the Delphix Management application.

Creating a Hook Operation Template

 The provisioning wizard still imports hook operations from templates.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Operation Templates**.
4. Click the **Plus** icon to add a new operation template.
5. Enter a **Name** for the template.
6. Select an operation **Type**.
7. Enter a **Description** detailing what the operation does or how to use it.
8. Enter operation **Contents** to implement the operation partially or fully.
9. Click **Create**.

Importing a Hook Operation Template

To import a hook operation template:

1. In the **Datasets** panel, select a dataset.
2. Click the **Configuration** tab.
3. Within the **Configuration** tab, click the **Hooks** tab.
4. Select the **hook** to edit.
5. Click the **Plus** icon to add a new operation.
6. Click **Import**.
7. Select the **template** to import.
8. Click **Import**.
9. When you have set all hook operations, click **Check** to save the changes.

Exporting a Hook Operation Template

To export a hook operation template:

1. In the **Datasets** panel, select a dataset.
2. Click the **Configuration** tab.
3. Within the **Configuration** tab, click the **Hooks** tab.
4. Select the **hook** to edit.
5. Click the **Plus** icon to add a new operation.
6. Select the **type of operation**.
7. Click the **text area** and edit the contents of the operation.
8. Click **Export**. 
9. Enter a **Name** for the template.

10. Enter a **Description** detailing what the operation does or how to use it.
11. Click **Export**.

Considerations for Delphix Hook Operations on Oracle RAC VDBs

When performing data management operations like SnapSync, refresh, and so on with Delphix data sources such as VDB and dSources, Delphix offers "hooks" as an ability to perform tasks prior to and after these data management operations. For an introduction and more details about hooks refer to, [Oracle Hooks for data sources](#).

Mostly the above article discusses shell operations within the hooks. But sometimes you might need to run a SQL script within the hook to capture or alter the database state before or after a Delphix operation on that source. In a sample scenario, we might be connecting to the database and running some SQL script that does the operations and perhaps generates some files for later use.

In such a case, there is no need to set the Oracle environment in the script, and in fact, setting the ORACLE_SID value in the script can be a problem for RAC instances. From the article linked above, these are the details of hook operations as they relate to RAC VDBs:

Oracle RAC

When linking from, or provisioning to Oracle RAC environments, hook operations will not run once on each node in the cluster. Instead, the Delphix Engine picks a node in the cluster at random and guarantees all operations within any single hook will execute serially on this node.

 The Delphix Engine does not guarantee the same node is chosen for the execution of every hook but does guarantee that Pre-/Post- hook pairs (such as Pre-Sync and Post-Sync) will execute on the same node.

If the data source (or VDB) is a single instance Oracle, explicitly setting variables in the script should not be a problem, since the local listener will be aware of the SID in the connect string. But for a RAC, since we will not be sure which node the hook will run against beforehand, do not statically allocate the variables, the engine will determine the connection string based on the databases section of the environments page of the admin app GUI.

Troubleshooting

Hook operations do not work for Delphix RAC installations.

Resolution

Remove any oracle environment setting parameters in the script and consider that the hook will run arbitrarily against any of the RAC nodes. Validate that the connect string works for the **Environment>Databases** section for the dSource in question.

Additional Information

For more information about validating and setting connection strings refer to, [Adding JDBC Connection String](#).

Related Topics

- [Hook Scripts for Automation and Customization](#)
- [Setting Hook Operations](#)

Cookbook of Common Scripts for Hooks on SQL Server

Customization is a huge topic and there are hundreds of ways to customize Delphix operations using hooks, but here are some example scripts which have shown up as common, and are presented here to provide guidance for building hooks.

- [Common Scripts For Hooks](#)
- [Example Powershell Script for Debugging](#)
- [Example PowerShell Script for Email Notification](#)
- [Example Powershell Script for Executing Stored Procedures](#)
- [Sample Configure Clone Hook for SQL Server](#)
- [Sample Post-Start Hook to Add CDC Jobs](#)
- [Sample Pre- and Post-Refresh Hook for SQL Server](#)

Common Scripts For Hooks

What Are Hooks?

Hooks are programmatic call-outs embedded before and after dataset manipulation operations within the Delphix Engine. These programmatic call-outs provide the ability for the Delphix Engine to run custom-built PowerShell scripts directly on the Windows database server.

Hooks and Data Operations

Delphix dSources are manipulated by "Sync" operations, in which the Delphix Engine captures data changes from backups of the source database appends them to the dSource. The pre-sync and post-sync hooks are executed in PowerShell by the "environment user" Windows domain account (delphix_trgt) on the staging target. Since the dSource is a database in recovery mode within the SQL instance on the staging target, no database operations can be performed on the dSource itself, but file-operations on the staging target are possible.

Delphix **VDBs** are manipulated by Provision, Refresh, Rewind, Snapshot, Start, and Stop operations. The Configure Clone, Pre-Refresh, Post-Refresh, Pre-Rewind, Post-Rewind, Pre-Snapshot, Post-Snapshot, Pre-Start, Post-Start, Pre-Stop, and Post-Stop hooks are executed in PowerShell by the Windows "environment user" Windows domain account (delphix_trgt) on the VDB on the target.

Here are some example use-cases and solutions using hooks:

- Your organization might require that a notification be sent to a distribution list five minutes before anyone refreshes or rewinds a VDB that is used for development or testing. You can write and test that logic in a PowerShell script and then embed it into the Pre-Refresh, Pre-Rewind, and Pre-Stop hooks on all VDBs. This PowerShell script can send email to a specified distribution list, and then pause for 300 seconds (5 minutes) before allowing the operation to continue.
 - [PowerShell scripting example for email notification during Pre-Refresh and Pre-Rewind hooks](#)
- Update data within a VDB when it is first provisioned and again after each refresh from the dSource. A very common example of this is data masking or anonymization, to prevent confidential data from being visible to the users and administrators of the VDB. Delphix has integrated data masking, but you can still initiate the appropriate job and poll until it completes using a PowerShell script.
- During a refresh of a VDB from the dSource, save the non-production database and application passwords, then restore them after the refresh is complete.
 - [PowerShell scripting example for stored procedure execution during Pre-Refresh and Post-Refresh hooks](#)
- Debug the sequence and timing of dataset manipulation operations within the Delphix Engine by logging hook executions to a file
 - [PowerShell scripting example for debugging output to be generated during all VDB hooks](#)

There are many more potential use-cases for hooks, but hopefully, these examples will begin to show the way.

Related Topics

- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Example Powershell Script for Debugging

Description

This is the Windows PowerShell code for a script named **display.ps1**, which you can insert into any pre-script, post-script, or hook to display a timestamp and the values of the Windows environment variables to a log file located in the **C:\TEMP** directory on the target host. One purpose of this script, if inserted into every hook, is to better understand the order of operations within the Delphix Engine. It would also help you understand the duration of those operations.

Of course, this basic script is intended to be augmented for more custom diagnostic output, as desired.

 **Disclaimer:** This sample code is offered for illustration purposes only, and does not carry any warranty or guarantee. Employing copies of this code, in whole or in part, indicates acceptance of all risks.

Source Code for Script

Downloadable copy [display.ps1](#).

```
#=====
# File:      display.ps1
# Type:      Powershell script
# Date:      04-Nov 2016
# Author:    Delphix
#
# Licensed under the Apache License, Version 2.0 (the "License");
# you may not use this file except in compliance with the License.
# You may obtain a copy of the License at
#
#   http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" BASIS,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
# See the License for the specific language governing permissions and
# limitations under the License.
#
# Copyright (c) 2016 by Delphix. All rights reserved.
#
# Description:
#
#   Windows Powershell script intended to be used for debugging purposes, by
#   logging timestamped variable values with the name of the hook from which
#   it was called to an output file on the user's desktop.
#
# Usage:
#
#   display.ps1 <hook-name>
#
#   where:
#       <hook-name>    the name of the dSource or VDB hook from
#                       which this script is called:
```

```

#           For dSources...
#           Pre-Sync, Post-Sync
#           For VDBs...
#           Configure-Clone,
#           Pre-Refresh, Post-Refresh,
#           Pre-Rewind, Post-Rewind,
#           Pre-Snapshot, Post-Snapshot
#
#   If the hook-name does not have the correct text, then the script will not
#   not recognize it and will emit an error message.
#
# Modifications:
#=====
param([string]$hookName)
#
#-----
# Set up variables for use within the script...
#-----
$outDir = "C:\TEMP"
$timeStamp = Get-Date -format o
#
#-----
# Output values for Windows environment variables set with hook session by the
# Delphix virtualization engine...
#
# If the name of the hook implies a VDB hook, then output the values for the
# relevant variables...
#-----
if ( ( $hookName.ToUpper() -eq "CONFIGURE-CLONE" ) -or
      ( $hookName.ToUpper() -eq "PRE-REFRESH" ) -or
      ( $hookName.ToUpper() -eq "POST-REFRESH" ) -or
      ( $hookName.ToUpper() -eq "PRE-REWIND" ) -or
      ( $hookName.ToUpper() -eq "POST-REWIND" ) -or
      ( $hookName.ToUpper() -eq "PRE-SNAPSHOT" ) -or
      ( $hookName.ToUpper() -eq "POST-SNAPSHOT" ) ) {
    $outFile = $outDir + "\HOOK_OUTPUT_" + $env:VDB_DATABASE_NAME + ".TXT"
    $timeStamp + "|" + $hookName + ": VDB_INSTANCE_HOST is '" +
$env:VDB_INSTANCE_HOST + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": VDB_INSTANCE_PORT is '" +
$env:VDB_INSTANCE_PORT + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": VDB_INSTANCE_NAME is '" +
$env:VDB_INSTANCE_NAME + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": VDB_DATABASE_NAME is '" +
$env:VDB_DATABASE_NAME + "'" | Out-File -FilePath $outFile -Append
} else {
    #
    #-----
    # If the name of the hook implies a dSource hook, then output the values
    # for the relevant variables...
    #-----
    if ( ( $hookName.ToUpper() -eq "PRE-SYNC" ) -or
          ( $hookName.ToUpper() -eq "POST-SYNC" ) ) {
        $outFile = $outDir + "\HOOK_OUTPUT_" + $env:SOURCE_DATABASE_NAME + ".TXT"

```

```

    $timeStamp + "|" + $hookName + ": SOURCE_INSTANCE_HOST is '" +
$env:SOURCE_INSTANCE_HOST + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": SOURCE_INSTANCE_PORT is '" +
$env:SOURCE_INSTANCE_PORT + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": SOURCE_INSTANCE_NAME is '" +
$env:SOURCE_INSTANCE_NAME + "'" | Out-File -FilePath $outFile -Append
    $timeStamp + "|" + $hookName + ": SOURCE_DATABASE_NAME is '" +
$env:SOURCE_DATABASE_NAME + "'" | Out-File -FilePath $outFile -Append
  } else {
    #
    #-----
    # If the name of the hook is not recognized, then output an error
    # message to that effect...
    #-----
    $outFile = $outDir + "\HOOK_OUTPUT_UNKNOWN.TXT"
    $timeStamp + ": unknown hook name '" + $hookName + "'" | Out-File -FilePath
$outFile -Append
  }
}

```

Output

The name of the output file created by this script includes the name of the database, so all output from all hooks associated with this database name will be appended to the same file.

Related Topics

- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Example PowerShell Script for Email Notification

Description

The PowerShell script (below) sends an email using SMTP attributes passed on the command-line (i.e. from-Email-Address, to-Email-Address, and SMTP-Server). The script then pauses for 5 minutes before returning control to the Delphix Engine.

Source Code for Script

Downloadable copy [email_5mins.ps1](#).

Description

The PowerShell script (below) sends an email using SMTP attributes passed on the command-line (i.e. from-Email-Address, to-Email-Address, and SMTP-Server). The script then pauses **for 5** minutes before returning control to the Delphix Engine.

Source Code **for** Script

Downloadable copy [email_5mins.ps1](#).

```

=====
# File:          email_5mins.ps1
# Type:          powershell script
# Author:       Delphix Professional Services
# Date:         04-Nov 2016
#
# Copyright and license:
#
#   Licensed under the Apache License, Version 2.0 (the "License"); you may
#   not use this file except in compliance with the License.
#
#   You may obtain a copy of the License at
#
#       http://www.apache.org/licenses/LICENSE-2.0
#
#   Unless required by applicable law or agreed to in writing, software
#   distributed under the License is distributed on an "AS IS" basis,
#   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
#
#   See the License for the specific language governing permissions and
#   limitations under the License.
#
#   Copyright (c) 2015,2016 by Delphix. All rights reserved.
#
# Description:
#
#   Script will send an email using the SMTP attributes passed as parameters
#   on the command-line, and then pause for 5 minutes.
#
#

```

```

# Command-line parameters:
#
#   fromEmailAddr  - originating email address (used for "reply-to")
#   toEmailAddr    - destination email address
#   smtpServer     - IP address or hostname of SMTP server
#   opName         - free-format text describing the Delphix operation
#                   about to occur (i.e. "REFRESH", "REWIND", "STOP", etc)
#
#
# Environment inputs expected:
#
#   VDB_DATABASE_NAME
#
# Note:
#
# Modifications:
#   TGorman   04nov16   first version
#=====
param(  [string]$fromEmailAddr = "~~~"
        ,[string]$toEmailAddr = "~~~"
        ,[string]$smtpServer = "~~~"
        ,[string]$opName = "~~~"
)
#
#-----
# Compose the email within a "Net.Mail.MailMessage" object, then send it...
#-----
$message = New-Object Net.Mail.MailMessage
$message.From = $fromEmailAddr
$message.To.Add($toEmailAddr)
$message.Subject = "VDB '" + $env:VDB_DATABASE_NAME + "' will be " + $opName + " in 5
mins"
$message.Body = "<H1>Delphix operation " + $opName + " will occur in 5 mins</H1>"
$message.Body = $message.Body + "<BR><BR>Please save any work in this database or it
will be lost."
$message.Body = $message.Body + "<BR><BR>Contact IT Operations at x5555 if you have
any questions or concerns"
$message.IsBodyHTML = "True"
$smtp = New-Object Net.Mail.SmtpClient($smtpServer)
$smtp.Send($message)
#
#-----
# Go to sleep for 300 seconds (5 minutes)...
#-----
Start-Sleep -s 300
exit 0

```

Related Topics
Cookbook of Common

Related Topics


- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Example Powershell Script for Executing Stored Procedures

Description

The Powershell script (below) makes a call to the stored procedure, whose fully-qualified name is passed as a parameter. This particular script is written so that the stored procedure expects only one parameter named "@DatabaseName". Please note that pre and post-scripts can only be executed on VDBs on the VDB target host. If "@DatabaseName" is not unique enough, then the script can be modified to add "@InstanceName" as well.

The script logs debugging information to a log file created within the "Desktop" directory of the Delphix "environment user" Windows account. This log file is directed to "C:\TEMP" in this example, but of course, can be modified as suitable in your environment.



Source Code for Script

Downloadable copy [callsp.ps1](#).

```

#=====
# File:      callsp.ps1
# Type:     powershell script
# Author:   Delphix Professional Services
# Date:    02-Nov 2015
#
# Copyright and license:
#
    
```

```

# Licensed under the Apache License, Version 2.0 (the "License"); you may
# not use this file except in compliance with the License.
#
# You may obtain a copy of the License at
#
#     http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" basis,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
#
# See the License for the specific language governing permissions and
# limitations under the License.
#
# Copyright (c) 2015 by Delphix. All rights reserved.
#
# Description:
#
# Call the appropriate stored procedure within the DBO schema in the MSDB
# database on behalf of the VDB. The stored procedure name the name of the
# database as a parameter called "@DatabaseName"..
#
# Command-line parameters:
#
# $fqSpName          fully-qualified stored procedure name
#
# Environment inputs expected:
#
# VDB_DATABASE_NAME  SQL Server database name for the VDB
# VDB_INSTANCE_NAME  SQL Server instance name for the VDB
# VDB_INSTANCE_PORT  SQL Server instance port number for the VDB
# VDB_INSTANCE_HOST  SQL Server instance hostname for the VDB
#
# Note:
#
# Modifications:
#   TGorman   02nov15   first version
#=====
param([string]$fqSpName = "~~~")
#
#-----
# Verify the "$dirPath" and "$fqSpName" command-line parameter values...
#-----
if ( $fqSpName -eq "~~~" ) {
    throw "Command-line parameter 'fqSpName' not found"
}
#
#-----
# Clean up a log file to capture future output from this script...
#-----
$dirPath = [Environment]::GetFolderPath("Desktop")
$timeStamp = Get-Date -UFormat "%Y%m%d_%H%M%S"

```

```

$logFile = $dirPath + "\" + $env:VDB_DATABASE_NAME + "_" + $timeStamp + "_SP.LOG"
"logFile is " + $logFile

#
#-----
# Output the variable names and values to the log file...
#-----
"INFO: dirPath = '" + $dirPath + "'" | Out-File $logFile
"INFO: fqSpName = '" + $fqSpName + "'" | Out-File $logFile -Append
"INFO: env:VDB_INSTANCE_HOST = '" + $env:VDB_INSTANCE_HOST + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_NAME = '" + $env:VDB_INSTANCE_NAME + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_PORT = '" + $env:VDB_INSTANCE_PORT + "'" | Out-File $logFile
-Append
"INFO: env:VDB_DATABASE_NAME = '" + $env:VDB_DATABASE_NAME + "'" | Out-File $logFile
-Append
#
#-----
# Housekeeping: remove any existing log files older than 15 days...
#-----
"INFO: removing log files older than 15 days..." | Out-File $logFile -Append
$ageLimit = (Get-Date).AddDays(-15)
$logFilePattern = $env:VDB_DATABASE_NAME + "*_SP.LOG"
"INFO: logFilePattern = '" + $logFilePattern + "'" | Out-File $logFile -Append
Get-ChildItem -Path $dirPath -recurse -include $logFilePattern |
    Where-Object { !$_.PSIsContainer -and $_.CreationTime -lt $ageLimit } |
    Remove-Item
#
#-----
# Run the stored procedure...
#-----
"INFO: Running stored procedure '" + $fqSpName + "' within database '" +
    $env:VDB_DATABASE_NAME + "'..." | Out-File $logFile -Append
try {
    "INFO: open SQL Server connection..." | Out-File $logFile -Append
    $sqlServer = $env:VDB_INSTANCE_HOST + "\" + $env:VDB_INSTANCE_NAME + ", " +
    $env:VDB_INSTANCE_PORT
    "INFO: sqlServer = '" + $sqlServer + "'" | Out-File $logFile -Append
    [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SqlServer.Smo") |
    Out-Null;
    $conn = New-Object System.Data.SqlClient.SqlConnection
    $conn.ConnectionString = "Server=$sqlServer; Database=MSDB; Integrated
Security=SSPI;"
    "INFO: conn.ConnectionString = '" + $conn.ConnectionString + "'" | Out-File
$logFile -Append
    $conn.Open()
    $cmd1 = New-Object System.Data.SqlClient.SqlCommand($fqSpName, $conn)
    $cmd1.CommandType = [System.Data.CommandType]::StoredProcedure
    $cmd1.Parameters.Add('@DatabaseName', $env:VDB_DATABASE_NAME) | Out-null
    "INFO: calling " + $fqSpName + ", @DatabaseName = " + $env:VDB_DATABASE_NAME |
    Out-File $logFile -Append
}

```

```
$exec1 = $cmd1.ExecuteReader()
$exec1.Close()
$conn.Close()
} catch { Throw $Error[0].Exception.Message | Out-File $logFile -Append }
#
"INFO: completed stored procedure '" + $fqSpName + "' within database '" +
  $env:VDB_DATABASE_NAME + "' successfully" | Out-File $logFile -Append
#
#-----
# Exit with success status...
#-----
exit 0
```

Related Topics

- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Sample Configure Clone Hook for SQL Server

How the Configure Clone Hook Works

The Delphix Engine executes the **Configure Clone** hook after the initial provisioning of a virtual database (VDB) and after each refresh of that VDB. However, during refresh operations, the **Configure Clone** hook runs before the **Post-Refresh** hook.

Because it executes after the Delphix Engine brings fresh data from the dSource into the VDB, the Configure Clone hook is an ideal place for data masking and other obfuscation and reconfiguration operations to occur.

The Windows PowerShell script (shown below) is intended to demonstrate how to call the HTTP interface of the Delphix DmSuite v4.7.3 from a Configure Clone hook.

 **Disclaimer:** This sample code is offered for illustration purposes only, and does not carry any warranty or guarantee. Employing copies of this code, in whole or in part, indicates acceptance of all risks.

```
#=====
# File:          mask.ps1
# Type:          powershell script
# Author:        Delphix Professional Services
# Date:          02-Nov 2015
#
# Copyright and license:
#
# Licensed under the Apache License, Version 2.0 (the "License"); you may
# not use this file except in compliance with the License.
#
# You may obtain a copy of the License at
#
# http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" basis,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
#
# See the License for the specific language governing permissions and
# limitations under the License.
#
# Copyright (c) 2015 by Delphix. All rights reserved.
#
# Description:
#
# Powershell script to automate the call to Delphix/AXIS data-masking when
# the appropriate parameters are passed.
#
# Otherwise, this script just returns without doing anything.
#
# Command-line parameters:
#
# $dmSuiteJobId          Delphix/AXIS DmSuite Job ID value
# $dmSuiteAppName        Delphix/AXIS DmSuite Application Name
```

```

# $dmSuiteEnvName          Delphix/AXIS DmSuite Target Environment Name
# $dmSuiteConnName        Delphix/AXIS DmSuite Target Connector Name
#
# Environment inputs expected:
#
# VDB_DATABASE_NAME       SQL Server database name for the VDB
# VDB_INSTANCE_NAME       SQL Server instance name for the VDB
# VDB_INSTANCE_PORT       SQL Server instance port number for the VDB
# VDB_INSTANCE_HOST       SQL Server instance hostname for the VDB
#
# Note:
#
# Modifications:
#=====
param([string]$dmSuiteJobId = "~~~"
      ,[string]$dmSuiteAppName = "~~~"
      ,[string]$dmSuiteEnvName = "~~~"
      ,[string]$dmSuiteConnName = "~~~"
)
#
#-----
# Set fixed variables for later use when submitting the Delphix/AXIS DmSuite job...
#-----
$dirPath = [Environment]::GetFolderPath("Desktop")
$dmSuiteServer = "XXXXXXXX"
$dmSuitePort = "8282"
$dmSuiteUser = "XXXXXXXX"
$dmSuitePwd = "XXXXXXXX"
$dmSuiteBaseURL = "http://" + $dmSuiteServer + ":" + $dmSuitePort + "/dmsuite/apiV4"
$dmSuiteWaitSecs = 5
#
#-----
# If no Delphix/AXIS DmSuite Job ID is provided, then this database is not to be
# masked, and so then don't do anything...
#-----
if ( $dmSuiteJobId -eq "~~~" -or $dmSuiteJobId -eq "0" ) {
    exit 0
}
#
#-----
# Otherwise, if we're going to mask this database, then first let's create a log
# file to capture output from this script and also create an XML file to store
# responses from the Delphix/AXIS DmSuite engine...
#-----
$timeStamp = Get-Date -UFormat "%Y%m%d_%H%M%S"
$logFile = $dirPath + "\" + $env:VDB_DATABASE_NAME + "_" + $timeStamp + "_MASK.LOG"
"logFile is " + $logFile
$xmlFile = $dirPath + "\" + $env:VDB_DATABASE_NAME + "_" + $timeStamp + "_MASK.XML"
"INFO: xmlFile = '" + $xmlFile + "'" | Out-File $logFile
$errFile = $dirPath + "\" + $env:VDB_DATABASE_NAME + "_" + $timeStamp + "_MASK.ERR"
"INFO: errFile = '" + $errFile + "'" | Out-File $logFile -Append
#
#-----

```

```

# Housekeeping: remove any existing log files older than 15 days...
#-----
"INFO: removing log files older than 15 days..." | Out-File $logFile -Append
$ageLimit = (Get-Date).AddDays(-15)
$logFilePattern = $env:VDB_DATABASE_NAME + "*_MASK.LOG"
"INFO: logFilePattern = '" + $logFilePattern + "'" | Out-File $logFile -Append
Get-ChildItem -Path $dirPath -recurse -include $logFilePattern |
    Where-Object { !$_.PSIsContainer -and $_.CreationTime -lt $ageLimit } |
    Remove-Item
#
#-----
# Record variables and parameters to the log file...
#-----
"INFO: dirPath = '" + $dirPath + "'" | Out-File $logFile -Append
"INFO: dmSuiteJobId = '" + $dmSuiteJobId + "'" | Out-File $logFile -Append
"INFO: dmSuiteAppName = '" + $dmSuiteAppName + "'" | Out-File $logFile -Append
"INFO: dmSuiteEnvName = '" + $dmSuiteEnvName + "'" | Out-File $logFile -Append
"INFO: dmSuiteConnName = '" + $dmSuiteConnName + "'" | Out-File $logFile -Append
"INFO: env:VDB_INSTANCE_HOST = '" + $env:VDB_INSTANCE_HOST + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_NAME = '" + $env:VDB_INSTANCE_NAME + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_PORT = '" + $env:VDB_INSTANCE_PORT + "'" | Out-File $logFile
-Append
"INFO: env:VDB_DATABASE_NAME = '" + $env:VDB_DATABASE_NAME + "'" | Out-File $logFile
-Append
"INFO: dmSuiteServer = '" + $dmSuiteServer + "'" | Out-File $logFile -Append
"INFO: dmSuitePort = '" + $dmSuitePort + "'" | Out-File $logFile -Append
"INFO: dmSuiteUser = '" + $dmSuiteUser + "'" | Out-File $logFile -Append
"INFO: dmSuiteBaseURL = '" + $dmSuiteBaseURL + "'" | Out-File $logFile -Append
"INFO: xmlFile = '" + $xmlFile + "'" | Out-File $logFile -Append
"INFO: errFile = '" + $errFile + "'" | Out-File $logFile -Append
#
#-----
# Verify that the CURL executable is available...
#-----
$curlExe = $dirPath.ToUpper() + "\CURL.EXE"
if (-Not (test-path $curlExe )) {
    "ERROR: Executable '" + $curlExe + "' not found; aborting..." | Out-File $logFile
-Append
    Throw "ERROR: Executable '" + $curlExe + "' not found; aborting..."
}
"INFO: curlExe = '" + $curlExe + "'" | Out-File $logFile -Append
#
#-----
# Obtain the CURL authentication token for the username and encrypted password...
#-----
"INFO: obtaining auth_token for username/encrypted-password..." | Out-File $logFile
-Append
$cmdLine = $curlExe + " -sIX GET `" + $dmSuiteBaseURL + "/login?user=" +
$dmSuiteUser + "&password=" + $dmSuitePwd + "`" 2>&1"
"INFO: powershell command-line is '" + $cmdLine + "'" | Out-File $logFile -Append
$output = invoke-expression $cmdLine

```

```

if ( $LastExitCode -ne 0) {
    "ERROR: GET auth_token failed; aborting..." | Out-File $logFile -Append
    $output | Out-File $logFile -Append
    Throw $output
}
if ($output | where {$_.ToUpper() -match "ERROR" -or $_.ToUpper() -match "FAILURE"})
{
    "ERROR: GET auth_token returned an error; aborting..." | Out-File $logFile
-Append
    $output | Out-File $logFile -Append
    Throw $output
}
$curlAuthTokenLine = $output | Select-String -pattern auth_token
$curlAuthToken = $curlAuthTokenLine -split '\s+' | Select-Object -Last 1
if ([string]::IsNullOrEmpty($curlAuthToken)) {
    "ERROR: NULL authentication token; aborting..." | Out-File $logFile -Append
    $output | Out-File $logFile -Append
    Throw $output
}
"INFO: auth_token '$' + $curlAuthToken + "' obtained" | Out-File $logFile -Append
#
#-----
# Retrieve the Delphix/AXIS DmSuite environment ID...
#-----
"INFO: retrieving DmSuite environment ID for dmSuiteAppName '$' + $dmSuiteAppName +
    '...' | Out-File $logFile -Append
$cmdLine = $curlExe + " -H `auth_token:" + $curlAuthToken + "`" + "`" + "`" +
    $dmSuiteBaseURL + "/environments`" > " + $xmlFile + " 2> " + $errFile
"INFO: powershell command-line is '$' + $cmdLine + "'" | Out-File $logFile -Append
$output = invoke-expression $cmdLine
if ( $LastExitCode -ne 0) {
    "ERROR: retrieving DmSuite environment ID failed; aborting..." | Out-File
$logFile -Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-Item $errFile
    Throw $output
}
if ($output | where {$_.ToUpper() -match "ERROR" -or $_.ToUpper() -match "FAILURE"})
{
    "ERROR: retrieving DmSuite environment ID returned an error; aborting..." | Out-
File $logFile -Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-Item $errFile
    Throw $output
}
$xml = Get-Content $xmlFile
$environmentURL = ""

```

```

$xml.SelectNodes("//Environment") | % {
    if (($_.Name -eq $dmSuiteEnvName) -and ($_.Application -eq $dmSuiteAppName)) {
        $environmentURL = $_.Link.href
    }
}
if ([string]::IsNullOrEmpty($environmentURL)) {
    "ERROR: retrieving DmSuite environment ID returned an error; aborting..." | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-Item $errFile
    Throw "ERROR: retrieving DmSuite environment ID returned an error; aborting..."
}
"INFO: environmentURL = '" + $environmentURL + "'" | Out-File $logFile -Append
Remove-item $xmlFile
Remove-Item $errFile
#
#-----
# Retrieve the Delphix/AXIS DmSuite connector ID...
#-----
"INFO: retrieving DmSuite connector ID for dmSuiteConnName '" + $dmSuiteConnName +
    "...'" | Out-File $logFile -Append
$cmdLine = $curlExe + " -H `auth_token:" + $curlAuthToken + "`" + "`" +
    $dmSuiteBaseURL + "/" + $environmentURL + "/connectors`" > " + $xmlFile + " 2> "
+ $errFile
"INFO: powershell command-line is '" + $cmdLine + "'" | Out-File $logFile -Append
$output = invoke-expression $cmdLine
if ( $LastExitCode -ne 0) {
    "ERROR: retrieving DmSuite connector ID failed; aborting..." | Out-File $logFile
-Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-item $errFile
    Throw $output
}
if ($output | where {$_.ToUpper() -match "ERROR" -or $_.ToUpper() -match "FAILURE"})
{
    "ERROR: retrieving DmSuite connector ID returned an error; aborting..." | Out-File $logFile -Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-item $errFile
    Throw $output
}
$xmlFile = Get-Content $xmlFile
$connectorURL = ""
$xml.SelectNodes("//Connector") | % {
    if ( $_.Name -eq $dmSuiteConnName ) {

```

```

        $connectorURL = $_.Link.href
    }
}
if ([string]::IsNullOrEmpty($connectorURL)) {
    "ERROR: retrieving DmSuite connector ID returned an error; aborting..." | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-item $errFile
    Throw "ERROR: retrieving DmSuite connector ID returned an error; aborting..."
}
"INFO: connectorURL = '" + $connectorURL + "'" | Out-File $logFile -Append
Remove-item $xmlFile
Remove-item $errFile
#
#-----
# Start the Delphix/AXIS DmSuitemasking job...
#-----
"INFO: starting DmSuite JobID " + $dmSuiteJobId + "..." | Out-File $logFile -Append
$cmdLine = $curlExe + " -X POST -H `\"auth_token:\" + $curlAuthToken +
    "`\" -H `\"Content-Type:application/xml`\" -d `\"<MaskingsRequest>\" +
###
### The two following lines are needed if the Delphix/AXIS DmSuite masking job is
"multitenant"...
###
###     "<Links>" + "<Link rel=`\"SourceConnector`\" href=`\"'" + $connectorURL + "`\"/>" +
###     "<Link rel=`\"TargetConnector`\" href=`\"'" + $connectorURL + "`\"/>" + "</Links>"
+
    "</MaskingsRequest>`\" `\"'" + $dmSuiteBaseUrl + "/applications/" + $dmSuiteAppName
+
    "/maskingjobs/" + $dmSuiteJobId + "/run`\" > " + $xmlFile + " 2> " + $errFile
"INFO: powershell command-line is '" + $cmdLine + "'" | Out-File $logFile -Append
$output = invoke-expression $cmdLine
if ( $LastExitCode -ne 0) {
    "ERROR: DmSuite job start returned failure exit code; aborting..." | Out-File
$logFile -Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-Item $errFile
    Throw "ERROR: DmSuite job start returned failure exit code; aborting..."
}
$xml = Get-Content $xmlFile
if ($xml.MaskingsResponse.ResponseStatus.Status -ne "SUCCESS") {
    "ERROR: DmSuite job run failed; aborting..." | Out-File $logFile -Append
    $output | Out-File $logFile -Append
    cat $xmlFile | Out-File $logFile -Append
    cat $errFile | Out-File $logFile -Append
    Remove-Item $xmlFile
    Remove-item $errFile
    Throw "ERROR: DmSuite job start failed; aborting..."
}

```

```

}
Remove-Item $xmlFile
Remove-item $errFile
#
#-----
# Check status of the submitted Delphix/AXIS DmSuite masking job...
#-----
$dmSuiteStatus = "RUNNING"
while ( $dmSuiteStatus -match "RUNNING" ) {
    #
    Start-Sleep -s $dmSuiteWaitSecs
    $totalWaitSecs = $totalWaitSecs + $dmSuiteWaitSecs
    #
    "INFO: Checking DmSuite JobID " + $dmSuiteJobId + "status after " +
        $totalWaitSecs + " secs..." | Out-File $logFile -Append
    $cmdLine = $curlExe + " -H `\"auth_token:\" + $curlAuthToken +
        "`\" -H `\"Content-Type:application/xml`\" `\" +
        $dmSuiteBaseUrl + "/applications/" + $dmSuiteAppName +
        "/maskingjobs/" + $dmSuiteJobId + "/results`\" > " +
        $xmlFile + " 2> " + $errFile
    "INFO: powershell command-line is '" + $cmdLine + "'" | Out-File $logFile -Append
    $output = invoke-expression $cmdLine
    if ( $LastExitCode -ne 0) {
        "ERROR: DmSuite job status returned failure exit status; aborting..." | Out-
File $logFile -Append
        $output | Out-File $logFile -Append
        cat $xmlFile | Out-File $logFile -Append
        cat $errFile | Out-File $logFile -Append
        Remove-Item $xmlFile
        Remove-Item $errFile
        Throw "ERROR: DmSuite job status returned failure exit status; aborting..."
    }
    [xml]$xml = Get-Content $xmlFile
    if ($xml.MaskingsResponse.ResponseStatus.Status -ne "SUCCESS") {
        "ERROR: DmSuite job status failed; aborting..." | Out-File $logFile -Append
        $output | Out-File $logFile -Append
        cat $xmlFile | Out-File $logFile -Append
        cat $errFile | Out-File $logFile -Append
        Remove-Item $xmlFile
        Remove-item $errFile
        Throw "ERROR: DmSuite job status failed; aborting..."
    }
    $dmSuiteStatus = $xml.MaskingsResponse.Maskings.Masking.Status
    #
}
#
#-----
# Return final exit status of completed Delphix/AXIS DmSuite masking job...
#-----
if ($dmSuiteStatus -eq "SUCCESS") {
    #
    "INFO: DmSuite JobID '" + $dmSuiteJobId + "' SUCCESS after " +
        $totalWaitSecs + " secs..." | Out-File $logFile -Append
}

```

```
Remove-Item $xmlFile
Remove-item $errFile
exit 0
#
} else {
#
cat $xmlFile | Out-File $logfile -Append
cat $errFile | Out-File $logfile -Append
"ERROR: DmSuite JobID '" + $dmSuiteJobId + "' failed with '" + $dmSuiteStatus +
" after " + $totalWaitSecs + " secs running..." | Out-File $logfile -Append
Throw "ERROR: DmSuite JobID '" + $dmSuiteJobId + "' failed with '" +
$dmSuiteStatus + " after " + $totalWaitSecs + " secs running..."
#
}
```

Related Topics

- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Sample Post-Start Hook to Add CDC Jobs

Sample post-start hook for CDC

This script will add Change Data Capture (CDC) capture and cleanup jobs on provisioned virtual databases.

```
$VDB_NAME = $env:VDB_DATABASE_NAME

$VDB_INSTANCE = $env:VDB_INSTANCE_NAME

SQLCMD -b -d $VDB_NAME -r0 -Q "EXEC sys.sp_cdc_add_job @job_type = N'capture';EXEC
sys.sp_cdc_add_job @job_type = N'cleanup'" -U <Username> -P <Password> -S ".\
$VDB_INSTANCE"

exit $LastExitCode
```

Sample post-start hook for CDC if provisioning is done from a lower database version to SQL2016 or above

This script will add CDC capture, cleanup jobs, and upgrade CDC metadata on provisioned virtual databases.

```
$VDB_NAME = $env:VDB_DATABASE_NAME

$VDB_INSTANCE = $env:VDB_INSTANCE_NAME

SQLCMD -b -d $VDB_NAME -r0 -Q "EXEC sys.sp_cdc_add_job @job_type = N'capture';EXEC
sys.sp_cdc_add_job @job_type = N'cleanup';EXEC sys.sp_cdc_vupgrade;" -U <Username> -P
<Password> -S ".\$VDB_INSTANCE"

exit $LastExitCode
```

 Place real username and password at respective placeholders while using the scripts.

Sample Pre- and Post-Refresh Hook for SQL Server

How Pre- and Post-Refresh Hooks Work

The Delphix Engine executes the Pre-Refresh hook before each refresh operation on a virtual database (VDB). It executes the Post-Refresh hook after each refresh operation. If a Configure Clone hook is also defined, then the Post-Refresh hook executes after all Configure Clone hooks.

Generally, after the initial provisioning of a VDB, changes are made to that VDB which should be saved across refreshes. Examples include resetting passwords to production accounts, adding non-production accounts, database references to other databases in production or non-production, and so on.

The Pre-Refresh hook is the first step of a process to capture the data to be saved, ending with the Post-Refresh hook to re-apply all that was captured, after the refresh has completed. So, the order of execution is:

1. Zero, one, or more **Pre-Refresh** hook(s)
2. The **Refresh** operation itself on the VDB
3. Zero, one, or more **Post-Refresh** hook(s)

The Windows PowerShell script (shown below) is intended to demonstrate how to call a SQL Server stored procedure, taking a single parameter for the VDB database name. In a Pre-Refresh hook, a custom-built stored procedure (named MSDB.DBO.CAPTURE_DB_SETTINGS) can be called to capture all data, and store it in (for example) the MSDB system database, or in a file on the Windows target host server. In a Post-Refresh hook, a custom-built stored procedure (named MSDB.DBO.APPLY_DB_SETTINGS) can be called to access all the data captured by the Pre-Refresh hook and re-apply it into the newly-refreshed VDB.

So, assuming that the Windows PowerShell script below is stored in a file named CALLSP.PS1 within a directory on the VDB target host named D:\DELPHIX\SCRIPTS, the call syntax within the Pre-Refresh hook might look something like this:

```
D:\Delphix\Scripts\callsp.ps1 MSDB.DBO.CAPTURE_DB_SETTINGS
```

...and the call syntax within the Post-Refresh hook might look something like this:

```
D:\Delphix\Scripts\callsp.ps1 MSDB.DBO.APPLY_DB_SETTINGS
```

Be aware that Delphix hooks set the following Windows environments from a VDB:

- VDB_DATABASE_NAME – the name of the Delphix VDB (not the SQL Server database)
- VDB_INSTANCE_NAME – the name of the SQL Server instance
- VDB_INSTANCE_PORT – the port number of the SQL Server instance
- VDB_INSTANCE_HOST – the name of the Windows host on which the SQL Server instance resides

You can access the values in these environment variables within PowerShell using the `$env:variable-name` syntax, as shown in the code below.

Sample code

Below is a sample code for the **callsp.ps1** script.

 **Disclaimer:** This sample code is offered for illustration purposes only, and does not carry any warranty or guarantee. Employing copies of this code, in whole or in part, indicates acceptance of all risks.

```
#=====
# File:          callsp.ps1
# Type:         powershell script
# Author:       Delphix Professional Services
```

```

# Date:          02-Nov 2015
#
# Copyright and license:
#
# Licensed under the Apache License, Version 2.0 (the "License"); you may
# not use this file except in compliance with the License.
#
# You may obtain a copy of the License at
#
#     http://www.apache.org/licenses/LICENSE-2.0
#
# Unless required by applicable law or agreed to in writing, software
# distributed under the License is distributed on an "AS IS" basis,
# WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
#
# See the License for the specific language governing permissions and
# limitations under the License.
#
# Copyright (c) 2015 by Delphix. All rights reserved.
#
# Description:
#
# Call the appropriate stored procedure within the DBO schema in the MSDB
# database on behalf of the VDB. The stored procedure name the name of the
# database as a parameter called "@DatabaseName"..
#
# Command-line parameters:
#
#   $fqSpName      fully-qualified stored procedure name
#
# Environment inputs expected:
#
#   VDB_DATABASE_NAME      SQL Server database name for the VDB
#   VDB_INSTANCE_NAME      SQL Server instance name for the VDB
#   VDB_INSTANCE_PORT      SQL Server instance port number for the VDB
#   VDB_INSTANCE_HOST      SQL Server instance hostname for the VDB
#
# Note:
#
# Modifications:
#=====
param([string]$fqSpName = "~~~")
#
#-----
# Verify the "$dirPath" and "$fqSpName" command-line parameter values...
#-----
if ( $fqSpName -eq "~~~" ) {
    throw "Command-line parameter 'fqSpName' not found"
}
#
#-----
# Clean up a log file to capture future output from this script...
#-----

```

```

$dirPath = [Environment]::GetFolderPath("Desktop")
$timeStamp = Get-Date -UFormat "%Y%m%d_%H%M%S"
$logFile = $dirPath + "\" + $env:VDB_DATABASE_NAME + "_" + $timeStamp + "_SP.LOG"
"logFile is " + $logFile
#
#-----
# Output the variable names and values to the log file...
#-----
"INFO: dirPath = '" + $dirPath + "'" | Out-File $logFile
"INFO: fqSpName = '" + $fqSpName + "'" | Out-File $logFile -Append
"INFO: env:VDB_INSTANCE_HOST = '" + $env:VDB_INSTANCE_HOST + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_NAME = '" + $env:VDB_INSTANCE_NAME + "'" | Out-File $logFile
-Append
"INFO: env:VDB_INSTANCE_PORT = '" + $env:VDB_INSTANCE_PORT + "'" | Out-File $logFile
-Append
"INFO: env:VDB_DATABASE_NAME = '" + $env:VDB_DATABASE_NAME + "'" | Out-File $logFile
-Append
#
#-----
# Housekeeping: remove any existing log files older than 15 days...
#-----
"INFO: removing log files older than 15 days..." | Out-File $logFile -Append
$ageLimit = (Get-Date).AddDays(-15)
$logFilePattern = $env:VDB_DATABASE_NAME + "*_SP.LOG"
"INFO: logFilePattern = '" + $logFilePattern + "'" | Out-File $logFile -Append
Get-ChildItem -Path $dirPath -recurse -include $logFilePattern |
    Where-Object { !$_.PSIsContainer -and $_.CreationTime -lt $ageLimit } |
    Remove-Item
#
#-----
# Run the stored procedure...
#-----
"INFO: Running stored procedure '" + $fqSpName + "' within database '" +
    $env:VDB_DATABASE_NAME + "'..." | Out-File $logFile -Append
try {
    "INFO: open SQL Server connection..." | Out-File $logFile -Append
    $sqlServer = $env:VDB_INSTANCE_HOST + "\" + $env:VDB_INSTANCE_NAME + ", " +
    $env:VDB_INSTANCE_PORT
    "INFO: sqlServer = '" + $sqlServer + "'" | Out-File $logFile -Append
    [System.Reflection.Assembly]::LoadWithPartialName("Microsoft.SqlServer.Smo") |
    Out-Null;
    $conn = New-Object System.Data.SqlClient.SqlConnection
    $conn.ConnectionString = "Server=$sqlServer; Database=MSDB; Integrated
Security=SSPI;"
    "INFO: conn.ConnectionString = '" + $conn.ConnectionString + "'" | Out-File
$logFile -Append
    $conn.Open()
    $cmd1 = New-Object System.Data.SqlClient.SqlCommand($fqSpName, $conn)
    $cmd1.CommandType = [System.Data.CommandType]::StoredProcedure
    $cmd1.Parameters.Add('@DatabaseName', $env:VDB_DATABASE_NAME) | Out-null
    "INFO: calling " + $fqSpName + ", @DatabaseName = " + $env:VDB_DATABASE_NAME |
    Out-File $logFile -Append

```

```
$exec1 = $cmd1.ExecuteReader()
$exec1.Close()
$conn.Close()
} catch { Throw $Error[0].Exception.Message | Out-File $logFile -Append }
#
"INFO: completed stored procedure '" + $fqSpName + "' within database '" +
  $env:VDB_DATABASE_NAME + "' successfully" | Out-File $logFile -Append
#
#-----
# Exit with success status...
#-----
exit 0
```

Related Topics

- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [Hooks for SQL Server](#)

Managing Policies

These topics describe creating and managing SnapSync, LogSync, Retention, and VDB Refresh policies.

- [Policies for Scheduled Jobs](#)
- [Creating Custom Policies](#)
- [Policies and Time Zones](#)
- [Configuring Retention on Individual Snapshots](#)

Policies for Scheduled Jobs

Introduction

Creating policies is a great way to automate the management of datasets in Delphix. Whether for syncing data, taking snapshots, or refreshing virtual databases, policies ensure that data is ready when needed.

In order to create policies, navigate to the 'Manage' dropdown and select 'Policies'.

There are five categories of policies that the Delphix Engine uses in conjunction with datasets objects:

- **SnapSync** – How often snapshots of a source database are taken for a dSource.
- **VDB Snapshot** – How often snapshots are taken of the virtual database (VDB).
- **Retention** – How long snapshots and log files are retained for dSources and VDBs.
- **VDB Refresh** – Automatic refresh of a VDB, either with the most recent snapshot or latest Timeflow logs. The default setting for this policy is **None**.
- **Replica Retention** – How long snapshots are retained on replicated namespaces for dSources and VDBs after they have been deleted on the replication source. Retention of snapshot applies as long as the database is not deleted on the source. If the database is deleted, then the next replication update will delete the database and the retained snapshots on the replication target.

Avoid Adding Conflicting Policies

The following example case can occur if the Snapshot policy is set to NONE, this has a direct impact on the retention policy and potential VDB data growth: An engine encountered a domain0 space issue that resulted from the archive logs on a single VDB taking up 25% of the pool. That engine was set to use a 'None' Snapshot policy along with a 1-week log retention policy. Because the archive logs were needed to provision from the one and only snapshot that existed, the engine was unable to enforce the retention policy.

There can be default or custom policies for each of these categories.

Policy Type	Description	User Access
Default	Default policies exist at the domain level and are applied across all objects in a category. The settings of a default policy in a category can be modified, but their names cannot be changed.	<ul style="list-style-type: none"> • Users with Delphix Admin credentials
Custom	Custom policies allow for the creation of unique policies to fit any schedule requirements. These can be setup with varying time intervals, ranging from minutes to days.	<ul style="list-style-type: none"> • Users with Delphix Admin credentials • Group and object owners

Setting Different Policies for Objects in a Group

Policies applied at the group level will affect all objects in that group. To set different policies for objects in a group, apply the policies at the group level first, then apply policies at the object level.

SnapSync Policies

SnapSync policies determine how often snapshots of the source database are taken.

In the Default SnapSync policy, a snapshot is taken daily at 3:30 AM local time and will cancel if not completed within four hours. If SnapSync does not complete within this four hour period, it will resume at the next scheduled daily time until the process is complete.

Click the 'Edit' icon to change the Default SnapSync policy, or click the 'Add' icon to create a new SnapSync policy.

Note

Policies may be configured using the provided Schedule date picker, by Interval, or by using a Quartz cron expression when selecting Custom.

Retention Policy

A Retention Policy defines how long the Delphix Engine retains snapshots and log files, which are used to rewind or provision objects from past points in time. The retention time for snapshots must be equal to or longer than the retention time for logs.

To support longer retention times, more storage may need to be allocated to the Delphix Engine. The retention policy – in combination with the SnapSync policy – can have a significant impact on the performance and storage consumption of the Delphix Engine. Retention policies can be customized to retain snapshots and logs for longer periods, enabling usage to specific points further back in time.

Replica Retention Policy

Replica Retention policy defines how long the snapshots are retained on replicated namespaces for dSources and VDBs, after they have been deleted on the replication source.

Normally, the snapshots that have been deleted on the replication source engine are also deleted on the replication target engine. A new retention policy is introduced to provide an extended lifetime of such snapshots on the replication target. The Replica Retention Policy is targeted and defined on the replication target. This policy can be applied to an entire replica namespace or could target specific groups or dSource/VDBs within the replica namespace.

Use the replica snapshots on the replication target to provision or refresh VDBs. Point-in-time provisioning may not be possible for these snapshots, this is done to optimize the disk space usage on the replication target.

 The replica retention policy can only be used to extend (**not to reduce**) the lifetime of the replica snapshots.

Extended replica snapshots can be deleted in the following cases:

- Similar to replicated snapshots, Extended replica snapshots cannot be deleted manually. They can only be removed by adjusting the policy duration so the snapshot is no longer covered.
- If the entire dSource or VDB is deleted, the extended replica snapshot will be deleted on the replication target as well.
- Oracle virtual PDBs require their CDB to be present and, if the CDB is deleted or is not replicated, then the extended replica snapshots on both the deleted CDB as well as the dependent virtual PDB will be deleted on the replication target. This can happen in a rare scenario where an Oracle virtual PDB is migrated from one CDB to another. Once the virtual PDBs snapshots that reference the old CDB are deleted, the old CDB can also be deleted. The snapshots on the virtual PDB that depend on the old CDB will not be covered by extended replica retention on the replication target.

The replica retention policy runs automatically on a schedule to cleanup expiring snapshots or whenever a replication receive job is executed.

Benefits of Longer Retention

With increased retention time for snapshots and logs, a longer (older) rollback period for data is allowed.

Common use cases for longer retention include:

- SOX compliance
- Frequent application changes and development
- Caution and controlled progression of data
- Reduction of project risk
- Speed of rollback or restoring to older points in time

Related Topics

- [Getting Started](#)

Creating Custom Policies

This topic describes creating custom policies based on cron expressions for specific database objects or groups.

Custom policies are created by editing a policy associated with a database object, either during its creation or through the **Policy Management** screen after it has been created. For information about creating custom policies for dSources and VDBs during the linking and provisioning processes, see the **Linking** and **Provisioning** topics listed for each data platform.

Note

Policies may be configured using the provided Schedule date picker, by Interval, or by using a Quartz cron expression by selecting Custom.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Policies**.
4. Select the **tab** for the object or group for which you want to create a policy.
5. In the selected tab click the **Add** button located on the right.
For example, to create a new VDB Snapshot policy:
 - a. Select the **VDB Snapshot** tab.
 - b. Click **+ VDB Snapshot**.
6. Enter **Name** for the policy.
7. Select a **Time Zone**, if this is not selected Delphix picks up the system default time.
8. Select **Timeout**. When a job is run if it does not finish in the selected time it will be killed.
9. Select a **Schedule**.
10. Select **Next**.
11. Select the object(s) to which you want to assign this policy.
12. Click **Submit**.

Policies and Time Zones

Policies and Time Zones

You can configure the SnapSync, VDB Snapshot, and VDB Refresh policies with the time zone in which the policy should be scheduled.

To edit the time zone of a policy:

1. Login to the **Delphix Management application**.
2. Click **Manage**.
3. Select **Policies**.
4. Click on the **pencil** icon located next to the selected policy.
5. Select the appropriate time zone from the drop-down list.
6. Click **Submit**.

 Retention and Quota policies are not schedulable and do not need a time zone.

Upgrading to Version 4.2 or higher:

Prior to version 4.2, a policy operated under the time zone of the policy's target. For example, a SnapSync Policy scheduled for 4:00 am every day that targeted a dSource in Eastern Standard Time (EST) and a dSource in Pacific Standard Time (PST) fired twice a day: once at 4:00 am EST and once at 4:00 am PST.

To maintain the same behavior of the Delphix Engine after an upgrade, the upgrade process clones existing policies with these clones differing only in their time zone. After upgrading, you may notice that the names of policies change to include the time zones in which they operate.

 Default policies are not cloned and always operate under the timezone of the Delphix Engine.

Example of an Upgrade Engine

In this example, the dSources and VDBs originally operated under either EST (America/New_York) or CST (America/Mexico_City), and new policies were created to reflect this.

Original Policy	New Policies
UserSnapSync	UserSnapSync (America/Mexico_City) UserSnapSync (America/New_York)
SnapshotTest	SnapshotTest (America/Mexico_City) SnapshotTest (America/New_York)
UserRefresh	UserRefresh (America/Mexico_City) UserRefresh (America/New_York)

After an upgrade, ensure that the policies are configured as expected; it may have been unclear prior to this upgrade when policies were actually firing.

Also, after upgrading to 4.2 or higher, you may consolidate/clean-up the clones and these changes will persist through future upgrades. If you go to the policy tab, and click on a policy you should see a timezone field. This

timezone field is editable. So for example, if you had "VDB_SNAP (US/Arizona)" and "VDB_SNAP (America/Phoenix)", you could delete one of the duplicates (they are both from the same time zone in this case), make sure the timezone field is set to the desired time zone and rename the remaining policy to "VDB_SNAP".

Configuring Retention on Individual Snapshots

This topic describes adding a custom retention definition for individual snapshots. This value will override that of the policy currently assigned to the container, for example, if 'forever' is selected then the snapshot will no longer be deleted via the retention policy.

Procedure

1. Log into the **Delphix Management** application.
2. Select **Resources > Storage Capacity**.
3. Expand the object (dsource or vdb) to modify.
4. Expand the snapshots. (it may take a few minutes for the individual snapshots to appear)
5. To configure the desired value in the 'keep until' column, click the figure in the column.
6. In the Keep For dialog either enter the **number of days** or tick **forever**.
7. Click **Save**.

Accessing the Delphix Engine

The admin and sysadmin User Roles

After installation, you will have two users: **sysadmin** and **admin**. Email addresses are required inputs for both accounts.

 When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

You can create additional sysadmin and admin users, as well as other users with more restricted privileges, as described in the topics under [Managing System Administrators](#).

User	Initial password	Abilities and Duties
sysadmin	sysadmin	<p>Typical system administration duties such as:</p> <ul style="list-style-type: none"> • modifying NTP, SNMP, SMTP settings • managing storage • downloading support logs for the Delphix Engine • performing upgrades and patches. <p>Launches the initial Server Setup configuration application</p> <p>Has access to the Command Line Interface (CLI)</p>
admin	delphix	<p>Manages data objects, all collectively referred to as the Delphix Engine "domain":</p> <ul style="list-style-type: none"> • dSources • virtual databases (VDBs) • users • groups • related policies and resources <p>Manages the Delphix Engine domain using either the browser-based Engine Admin application or the Command Line Interface (CLI).</p>

Access Methods

There are three main interfaces to a Delphix Engine:

- Web browser (http or https)
- Command Line Interface (CLI) through ssh
- Delphix Web APIs

Accessing a Delphix Engine with a Web Browser

In order to reach the main Engine Administrator and Delphix Management application (GUI), you can use http or https:

http://<engine IP or hostname>

or

https://<engine IP or hostname>

 In order to use your hostname, your DNS Administrator must add the hostname and IP Address to your DNS system.

The Delphix Management Application

By default, the GUI login screen of a Delphix Engine will log you into the Delphix Management application, as that is where a majority of the day-to-day activities will take place, such as attaching of source databases and VDB provisioning. There is a link (Setup) to the Delphix Server Setup application.

If you are already logged into the Delphix Setup application, then you can switch to the Delphix Management login by clicking “Management” in the upper-right corner of the Delphix Setup screen:

The Delphix Setup GUI

If accessing the Delphix Engine to perform system administrator actions (adding storage, changing system-level parameters), you will need to switch the login screen to the Delphix Setup login by clicking **Server Setup** on the default login screen.

If you are already logged into the Delphix Management GUI, you can switch to the Server Setup login screen by clicking the username in the upper-right corner of the GUI and selecting “Setup”:

Accessing a Delphix Engine using the CLI

Using the Delphix CLI is useful for various situations:

- Limited bandwidth connections where a GUI is too “heavy”
- Flash is prohibited on browsers
- Users who prefer not to use a GUI
- Automation scripts

You access the Delphix CLI through an SSH session. Like the GUI, there are separate CLIs for the Engine Administrator and Delphix Sysadmin roles. The appropriate CLI is automatically selected based on the role that has been granted to you.

Using your SSH client of choice, access the Delphix Engine in a method similar to the following:

```
ssh <user>@<engine IP or hostname>
```

Where <user> is a pre-defined user on the Delphix Engine, and <engine IP> is the IP address of a Delphix Engine.

```
ssh delphix_admin@dlpxengine1
```

All Delphix functionality is available via the CLI. For further information on using the Delphix CLI, see [Command Line Interface Guide](#).

Accessing a Delphix Engine Using the Delphix Web API

Delphix offers a set of RESTful web service APIs with which you can administer a Delphix Engine. Using Web API calls allows you to create powerful, complex automation, often in coordination with other technologies like Jenkins, Puppet, Chef, and Ansible. All Delphix functionality is available via the Web API. The use of the Web API is outside the scope of this document, but there is additional information in the [Delphix Web Service API Guide](#).

Related Links

- [Managing Policies](#)
- [Users and Groups](#)

Oracle Environments and Data Sources

This section covers the following topics:

- [Delphix Architecture with Oracle](#)
- [Transparent Data Encryption and Delphix](#)
- [Oracle Environment Management](#)
- [Quick Start Guide for Oracle on Linux](#)
- [Quick Start Guide for Oracle on Solaris SPARC](#)
- [Oracle Support and Requirements](#)
- [Managing Oracle Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with Oracle](#)
- [Provisioning and Managing Virtual Databases with Oracle](#)
- [Oracle Hook Operations](#)

Delphix Architecture with Oracle

Overview

Various Oracle configurations ranging from Oracle RAC to Oracle multi-tenant can be used with Delphix. This article contains an overview of how Delphix works with Oracle.

There are three key concepts when using Delphix with any data platform:

1. **Environments:** The server and software required to run a data set. For Oracle, this will be an operating system host with Oracle instances running on it.
1. **Source Environment:** Source data to be ingested into Delphix. These will be used to create dSources.
2. **Target Environment:** Target hosts to provision VDBs. These need Oracle installations that correspond to the versions of the Source environments, per our [Oracle Support Matrix](#).
 - **dSources:** A database that the Delphix Virtualization Engine uses to create and update virtual copies of your database
 - **VDBs:** A database provisioned from either a dSource or another VDB which is a copy of the source data. A VDB is created and managed by the Delphix Virtualization Engine.

Explore how Delphix connects to Oracle environments and creates Oracle dSources and VDBs, with these concepts in mind.

Delphix is not a replacement for an archived log backup solution and should not be relied on to do so. This can cause issues where Delphix is unable to start a vPDB that was plugged into a physical CDB if Delphix is having to determine and deliver the archived logs. You will need to find a separate archived log backup solution workflow to own and be responsible for physical CDB archived logs.

Environment Linking and Provisioning Architecture

As shown in the diagram below Delphix begins by ingesting data from your source databases in order to create dSources. Once you have added an environment, Delphix will automatically ‘discover’ databases on it. This means we are looking for any compatible data sources we could ingest from to create new dSources.



When you specify a dSource to add, we will leverage RMAN and JDBC in order to create snapshots of the source. We continuously monitor the source database for new backups based on your SnapSync and LogSync policies. The result is a TimeFlow with consistent points from which you can provision a VDB, and a faster provisioning process, because there is no need for any database recovery during provisioning.

When you later provision a VDB, you can specify any environment as a target, including the environment that contains the staging database. However, for best performance, we recommend that you choose a different target environment. It must have an operating system that is compatible with the one running on the source environment.

Related Topics

- [Oracle Support and Requirements](#)
- [Linking Data Sources and Syncing Data with Oracle](#)
- [Provisioning and Managing Virtual Databases with Oracle](#)


Transparent Data Encryption and Delphix

Overview

The Oracle Transparent Data Encryption feature encrypts the sensitive data (database tables and tablespaces) stored on the disk. This prevents misuse of the data if the disks or storage mediums are lost or stolen. The data is transparently decrypted for the authorized users when they access the data.


Data is encrypted with the help of encryption keys which are stored in an external module or file, known as the *wallet* or *keystore*. The keystore is managed by an authorized user and can be either a hardware or software keystore. In order to decrypt the data successfully, the encryption keys must be made available to the database by the authorized user.

In the Oracle database, the data is organized into tables that are located within a tablespace, which is in turn made up of one or more files on disk. With TDE either individual tables or an entire tablespace can be encrypted. The keys used to encrypt the data are stored within the database itself, in the data dictionary. The keys themselves are further encrypted using the wallet keys. By using 2 layers of encryption in this manner, access to all of the encrypted data can be achieved with just access to the keys in the keystore.




In the remainder of this document, the term *key* will refer to the master key stored in the external keystore, rather than the table keys stored in the data dictionary. In the case of a multi-tenant database, the same principle applies, namely that the keys are stored in the keystore which is located external to the database. However, not all of the PDBs within the container need to be encrypted. The CDB itself must always have a key associated with it, even if the CDB datafiles themselves are not encrypted. This is the case because Oracle will encrypt both the datafiles belonging to the PDB, as well as the archive logs, which are located in the CDB, as illustrated in the following diagram.

In the diagram below, PDBs 1 and 2 are encrypted, but PDB 3 is not. Note that the keystore contains all of the keys for the entire multitenant database, in this case, keys for the CDB, PDB 1, and PDB 2.



In order to unplug a TDE-encrypted PDB, the key(s) for that PDB must be exported from the keystore into a new file, known as an *exported keyfile*. Oracle will prevent the PDB from being unplugged unless the keys have first been exported. The exported keyfile is encrypted with a password (known as the *keyfile secret*) that is specified during the export. The diagram below illustrates the scenario where PDB 2 has been unplugged. Note that the key for PDB 2 still remains in the keystore. This illustrates an important point about Oracle keystores - once added, a key cannot be easily removed from a keystore.



You can plug back a PDB into a new CDB before importing the new keys. This results in the PDB being plugged in in restricted mode with plugin violations. Importing the keys is necessary to resolve the plugin violations and start up the PDB in unrestricted READ WRITE mode.

The design of TDE allows for the keys for a given CDB or PDB to be changed, via the `ADMINISTER KEY MANAGEMENT SET KEY` command. This process is known as *key rotation* and will update the master encryption key in the keystore while leaving the table keys in the data dictionary. The table keys themselves will be encrypted with the new master encryption key, and any future updates to that data will use the new key. Existing data already written to the datafile or archive log will still be encrypted with the original key, however. For this reason, keys are never removed from an existing wallet, only new keys are added. In order to decrypt all of the data in a given database, all current and prior keys will be needed.

For more information on TDE and how it is configured and enabled, see the [Oracle documentation](#).

Delphix Implementation of TDE

Oracle recommends that the keystore be stored on a separate disk from the datafiles. In accordance with this recommendation, neither keystores nor exported keyfiles are stored on Delphix storage. Rather, they are placed on customer storage. Exported keyfiles generated by Delphix are stored in the *artifact directory* under the *keystores root directory*, while keystores generated by Delphix are stored in the location specified by `sqlnet.ora` for the target container database. It is the customer's responsibility to maintain these storage locations and ensure they are backed up as needed, just like database files. If the keystore or exported keyfile is lost, the data in the associated vPDB may not be recoverable and the vPDB will cease to operate.

Delphix support for TDE has the following requirements:

- TDE-enabled vPDBs must be provisioned to a linked CDB, not a vCDB.
- The Oracle version must be 12cR2 or higher (12cR1 is not supported).
- System tablespaces (either in the CDB or PDB) must not be encrypted.
- Keystores must not be on ASM storage.
- Migration of a TDE-enabled vPDB from one CDB to another requires [manual steps](#) that must be completed for migrate vDB to be successful and to support refresh and rewind operations on the migrated vPDB.
- Only software keystores on the same host as the database files are supported, in particular, Oracle Key Vault is not supported.
- The dSource from which the initial provision is done must be encrypted when it is linked. Existing dSources can not be encrypted without unlinking and creating a new dSource.
- There is currently no supported transition path from existing TDE-enabled vPDBs using the TDE workaround to the full product solution. The TDE workaround continues to be supported for approved customers.
- The target container database should have an autologin keystore configured. For a cluster target, the autologin keystore is shared.

Some or all of these restrictions may be relaxed in future versions of Delphix.

Definitions

The following terms are used throughout TDE documentation and are summarized here for clarity. Note that the first occurrence of these terms may be on other documentation pages.

Term	Definition
Keystore/wallet	File found on the Oracle host which stores the keys used to encrypt and decrypt the internal table keys in a database. Every keystore has a password which is set when it is first created, and must be supplied for operations on it.

Parent keystore	Keystore with the keys used to encrypt the dSource PDB files.
Target keystore	Keystore for the target CDB into which the TDE-encrypted vPDB is plugged.
Artifact directory	Directory on the target system (not on Delphix storage) which stores keys needed to support Delphix workflows on TDE-enabled vPDBs. It is located under the keystores root directory.
Exported keyfile	File located on the target Oracle host which contains keys that have been exported from the keystore. It is encrypted with a secret that is specified when it is exported. The exported keyfile itself cannot be used as a keystore, but its contents can be imported into a new keystore.
Key rotation	Process for changing the master encryption key in the keystore via ADMINISTER KEY MANAGEMENT SET KEY. This does not remove the original key, rather it adds a new key to the wallet and future data will be encrypted with the new key.
Auxiliary container database (CDB)	Provisioning an Oracle vPDB requires running recovery to bring the snapshotted datafiles into a consistent state. This needs to be done in the context of a container database, which is created on the target system. After recovery is complete, the vPDB is unplugged and plugged into the target container, and the auxiliary container is deleted.
Keyfile secret	Password used to encrypt an exported keyfile.
Keystores root directory	User-specifiable location on the target system under which all TDE related artifacts such as keystores and exported keyfiles created by Delphix are stored. This includes both the artifact directories used for vPDBs and temporary directories used for auxiliary CDB keystores.

Wallet Location Configuration

Oracle requires that the keystore location be specified to the database so that it can be accessed when reading from or writing to the database files. This location can be specified in 2 ways:

1. The ENCRYPTION_WALLET_LOCATION parameter in `sqlnet.ora`.
2. The wallet_root initialization parameter is available starting in Oracle 18, while ENCRYPTION_WALLET_LOCATION is available in Oracle 12.2. Oracle has indicated that support for ENCRYPTION_WALLET_LOCATION will end starting with Oracle 20.

Delphix supports both configurations for the appropriate releases (i.e. ENCRYPTION_WALLET_LOCATION only in 12.2, and both ENCRYPTION_WALLET_LOCATION and wallet_root in 18+). When using ENCRYPTION_WALLET_LOCATION, Delphix recommends referencing an environment variable, for example:

Encryption Wallet Location

```
ENCRYPTION_WALLET_LOCATION=
(SOURCE=
(METHOD=FILE)
(METHOD_DATA=
(DIRECTORY=/u03/app/ora12201/admin/$ORACLE_UNQNAME/wallet/)))
```

As there is only one `sqlnet.ora` file found under `$ORACLE_HOME`, it will be used for all databases that use that home. Specifying an environment variable such as `$ORACLE_UNQNAME` allows a different location for each database in the same `$ORACLE_HOME`. Any environment variable referenced in `sqlnet.ora` must always be set in the environment for the Oracle user. Delphix explicitly sets only `$ORACLE_HOME`, `$ORACLE_SID`, and `$ORACLE_UNQNAME` in the connections which are established by the Delphix engine, so it is recommended that only these variables be referenced in `sqlnet.ora`. For a 12.2 TDE vPDB provision, Delphix creates a unique `sqlnet.ora` file for the use of the auxiliary database during the provision.

Required O/S Permissions for the Delphix User

The provision itself is executed within the context of the environment user specified during the provision. This user does not have to be the Oracle user, and in fact, often is not. The Delphix user must be a member of the oracle group. During a TDE-enabled vPDB provision, the parent keystore is copied from a user-specified location to a location under the keystores root directory. The Delphix user does this copy via `ADMINISTER SQL MANAGEMENT` command. Since the Oracle user will do this, the Oracle user must be able to also create files in the wallet location.

The privilege requirements are satisfied by ensuring that the parent keystore has group read privileges, and the keystores root directory (owned by the Delphix user) has group write privileges.

Related Topics

- [Provisioning a TDE-enabled vPDB](#)
- [Migrating a TDE-enabled vPDB](#)
- [Provisioning a vPDB from a Non-Multitenant Source](#)
- [Provisioning a TDE-enabled vPDB to a Cluster Target](#)
- [Refreshing and Rewinding a TDE-enabled vPDB](#)

Oracle Environment Management

In order for the Delphix Engine to ingest data or provision virtual databases, your Oracle environments must comply with Delphix environment configuration standards.

Related Topics

- [Oracle Environments and Data Sources](#)

Quick Start Guide for Oracle on Linux


Overview

This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with Oracle database objects in the Delphix Engine. It does not cover advanced configuration options including Oracle RAC, Linking to Standby, or best practices for performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix Engine functionality. It assumes you will use the VMware Hypervisor.

In this guide, we will walk through deploying a Delphix Engine, configuring Oracle Source and Target environments on Linux servers, creating a dSource, and provisioning a VDB.

The following diagram describes the engine topology for Oracle environments. It illustrates the recommended ports to be open from the engine to remote services, to the engine, and to the Source and Target Environments.

For this Quick Start Guide, you can ignore the following components: Validated Sync Host, Replicated Engine, and Delphix Masking Engine.



Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online.
This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
the primary interface. Setting this value
                    to "true" will cause all other properties (address,
hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
"1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
-- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24

```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

- Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`)

- Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

- Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the browser at:
```

```
http://<address>
```

Type "`exit`" to disconnect, or any other commands to **continue** using the CLI.


- Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
- Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Once you setup the network access for Delphix Engine, navigate to the Delphix Engine URL in your browser for server setup.

The welcome screen below will appear for you to begin your Delphix Engine setup.



The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at <http://<Delphix Engine>/login/index.html#serverSetup>.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

⚠ The default domain user created on Delphix Engines from 5.3.1 is known as **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications

Choose your option to setup system time in this section. For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click **Next** to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication Service

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.

4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

⚠ Although your Delphix Engine will work without registration, we strongly recommend that you register each engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

Requirements for Oracle Hosts and Databases

In order to begin using Oracle environments with Delphix, you will need to configure the source and target hosts with the requirements described on this page.

Oracle Hosts and Databases

On each host with Oracle, there must be an operating system user configured to the required specifications for Delphix, as explained in the table below. This user (i.e. delphix_os) can easily be created by using the createDelphixOSUser.sh script (located at the bottom of the page).

These requirements apply to both source and target environments. However, target environments have additional requirements which are detailed in the ‘Target Host Requirements’ section below.

Source Host Requirement

Host Requirement	Explanation
Profile and privileges should be the same as the Oracle user (e.g. oracle) on the host.	For example, delphix_os should have the same environment variable settings (\$PATH, \$ORACLE_HOME, etc.), umask, and ulimit settings, as the user oracle.

<p>The Delphix software owner account (e.g. delphix_os) must have the same primary OS group as the Oracle software owner account (e.g. oracle).</p>	<p>Delphix recommends giving the delphix_os user the same primary OS group as the Oracle home owner. This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.</p> <p>Often, this is an OS group named oinstall. However, the oinstall group is not always necessary depending on your Oracle configuration. This user requires access to the libobk_proxy.so library in the toolkit for the child processes triggered by RMAN.</p>
<p>The delphix_os user must have the Oracle OSDBA group (typically dba) as a primary or secondary OS group.</p>	<p>The OSDBA group allows "OS authentication" when connecting to an Oracle database instance by specifying either username or password (i.e. rman target /), thus eliminating the need to store or retrieve a SYSDBA password.</p> <p>Oracle 12c For Oracle 12c and later versions of Oracle, the delphix_os user can also use OSBACKUPDBA as its secondary group. This is typically the BACKUPDBA group on the host.</p>
<p>For secondary group requirements, the Delphix OS user must be a member of the SYSBACKUP OS group (12.1 or higher) or the SYSDBA OS group (11.2 and lower).</p>	<p>This ensures that the Delphix engine is able to take snapshots of source databases using RMAN.</p>
<p>There must be a directory on the source host where the Delphix Engine Toolkit can be installed, for example: /var/opt/delphix/toolkit</p>	<p>The delphix_os user and primary OS group (i.e. oinstall) must own the directory.</p> <p>The directory must have at least permissions -rwxrwx--- (0770)</p> <p>The delphix_os user must have read and execute permissions on each directory in the path leading to the toolkit directory.</p> <p>At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.</p>
<p>The Delphix Engine must be able to make an SSH connection to the source host (typically port 22).</p>	
<p>Read access to \$ORACLE_HOME and all underlying files and directories.</p>	<p>Delphix needs to run locally available Oracle tools such as sqlplus and RMAN. Those executables, as well as various required libraries, reside inside the Oracle home</p>

Additional requirements for RAC Environments

If the source host is part of a RAC cluster, Delphix will attempt to use all nodes and crsctl for its operations.

RAC Environment Requirement	Explanation
The delphix_os user and Delphix Toolkit configuration must be the same on each node in the RAC cluster	<ul style="list-style-type: none"> delphix_os must have execute permission on crsctl and srvctl on each node in the cluster. access to olsnodes is needed. The Delphix Toolkit must be installed in the same directory path on each of the nodes in the source cluster. All data files, archive logs, and database control file must be located on storage shared by all of the cluster nodes. Each node in the cluster must be able to access archive logs and the database control file from all other nodes. At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.

Note: Masked Provisioning is not supported on Oracle RAC.

Auto-Discovery Requirements

The preferred way to find source databases is to allow Delphix to automatically discover your Oracle Homes and Databases by examining the inventory and oratab files and the listener setup. Successful autodiscovery requires read access to these and related files.

In **most** environments, delphix_os group membership is sufficient to perform auto-discovery.

However, if you have overridden Oracle's group permission structure, you may need to modify privileges to allow auto-discovery.

Auto-Discovery Requirements	Explanation
<p>The ORATAB file must exist (typically in /etc/oratab or /var/opt/oracle/oratab) and be readable by delphix_os.</p> <p>Read access to:</p> <ul style="list-style-type: none"> /etc/orainst.loc or /var/opt/oracle/orainst.loc. the Oracle inventory file (inventory.xml) 	<p>The ORATAB is used to determine the location of the Oracle installations on your host.</p> <p>Note: The Oracle inventory file is identified by the contents of orainst.loc (for example, \$INVENTORY_HOME/ContentsXML/inventory.xml).</p>
<p>Permission to run pargs on Solaris hosts and ps on AIX, HP-UX, and Linux hosts, as super-user.</p> <p>This permission is usually granted via sudo authorization of the commands.</p>	<p>Unless you have used a custom TNS_ADMIN setting, elevated access to ps (pargs on Solaris) is not required</p> <p>See the topic Sudo Privilege Requirements for Oracle Environments for further explanation of this requirement, and Sudo File Configuration Examples for Oracle Environments for examples of file configurations</p>

Source Database Requirements

For each source database, there are specific configurations required for Delphix to ingest data.

Source Database Requirement	Explanation
For each Oracle Home, the delphix_os user should have execute permission for the programs in \$ORACLE_HOME/bin.	<p>If symlinks are configured, Delphix must be configured with the same \$ORACLE_HOME path as was used when starting the instance.</p> <p>Ensure the PermitUserEnvironment configuration parameter = "yes" in the sshd_config file</p> <p>The \$ORACLE_HOME/bin/oracle executable must have the SETUID and SETGID flags set. Permissions on the oracle binary must be at least -rwsr-s-x (06751).</p>
Source databases must be in ARCHIVELOG mode to ensure that redo logs are archived.	Archive logs are required to make SnapSyncs consistent and provisionable.
Enable Block Change Tracking (BCT). (Highly Recommended).	Enabling BCT will improve SnapSync operation time. Without BCT, incremental SnapSyncs must scan the entire database.
Enable FORCE LOGGING. (Highly Recommended)	<p>If you do not enable FORCE LOGGING and NOLOGGING operations take place, you will get a Fault from the Delphix Engine.</p> <p>If you must use NOLOGGING to meet specific performance criteria, take a new snapshot of the source database after doing the NOLOGGING operations to bring the dSource up-to-date before provisioning VDBs.</p>
If the online redo log files are located on RAW or ASM devices, then LogSync can only operate in Archive Only mode.	See the topics Advanced Data Management Settings for Oracle dSources and Linking Oracle Physical Standby Databases for more information.

Operating System Specific Requirements

Solaris

On a Solaris host, gtar must be installed. Delphix uses gtar to handle long file names when extracting the toolkit files into the toolkit directory on a Solaris host. The gtar binary should be installed in one of the following directories:

- [/bin/usr](#)
- [/bin/sbin/usr](#)
- [/sbin/usr/contrib](#)
- [/bin/usr/sfw](#)
- [/bin/opt/sfw](#)
- [/bin/opt/csw/bin](#)

Additional Target Host Requirements

This topic describes the user privileges and environment requirements that are required for Oracle target hosts and databases collectively referred to as target environments.

These are in addition to the ‘Oracle Hosts and Database Requirements’ called out above.

Target Host Requirement	Explanation
The Delphix OS user must be a member of the SYSDBA OS group.	This ensures that the Delphix engine is able to create new VDBs on target hosts.
There must be a directory (e.g. "/mnt/provision/") that will be used as a container for the NFS mount points that are created when provisioning a VDB.	<p>The delphix_os user and primary OS group (i.e. oinstall or dba) must own the directory.</p> <p>The directory must have at least permissions -rwxrwx--- (0770).</p> <p>There must be no symbolic links in the path of this directory, because NFS can mount into a directory with symlinks in its path, but cannot unmount.</p>
Permission to run: <ul style="list-style-type: none"> • mount, umount, mkdir, rmdir as super-user. • pargs on Solaris hosts • ps on Linux, AIX, and HP-UX as super-user. • nfso on AIX hosts as super-user 	The following permissions are usually granted via sudo authorization of the commands. See Sudo Privilege Requirements for Oracle Environments for further explanation of the commands, and Sudo File Configuration Examples for Oracle Environments for examples of the /etc/sudoers file on different operating systems.
(Optional) Write permission to the \$ORACLE_BASE_CONFIG/dbs directory.	(i.e. chmod g+w \$ORACLE_BASE_CONFIG/dbs)
An Oracle listener process running on the target host.	The listener's version should be equal to or greater than the highest Oracle version that will be used to provision a VDB.
Required packages on target hosts: <ol style="list-style-type: none"> 1. portmapper / rpcbind 2. status daemon (rpc.statd) 3. NFS lock manager (rpc.lockd/lockmgr) 	As the Delphix Engine uses NFSv3 for mounting target host filesystems, the prerequisite packages to support NFSv3 client communication are required for normal operation, and the required services to support NFS client communications (including file locking) must be running. NFSv3 is enabled by default and to enable NFSv4, see NFSv4 Configuration .
If using ASM, the delphix_os user must have the Oracle ASM ownership groups as secondary OS groups (typically asmadmin and asmdba)	This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.

Deploy Hostchecker to Validate Delphix Requirements

Delphix has developed a hostchecker script that contains standardized checks for source and target hosts - these checks generally fall into three buckets

- OS and Host permissions/access
- Network Port Checks
- DB-specific functionality

OS and Host permissions/access and network port checks can (and should) be performed prior to Delphix installation to ensure a smooth deployment.

Each DB should have a specifically associated hostchecker - you can find detailed documentation on the DB-specific hostchecker page.

1. Download the appropriate **HostChecker tarball** for your engine from <https://download.delphix.com/>. Tarballs follow the naming convention "hostchecker_<OS>_<processor>.tar". For example, if you are validating a linux x86 host you should download the hostchecker_linux_x86.tar tarball.
2. Create a working directory and extract the **HostChecker files** from the **HostChecker tarball**.

```
mkdir dlp-host-checker
cd dlp-host-checker/
tar -xf hostchecker_linux_x86.tar
```

3. Run the `sh` script contained within:

```
sh hostchecker.sh
```

This will extract the JDK included in the tarball (if necessary) and invoke the HostChecker.

```
ora10205@bbdhcp:/home/ora10205/hostchecker-> sh hostchecker.sh
Extracting the JDK from the tarball jdk-6u45-linux-i586.tar.gz.
```

Don't Run as Root

Do not run the HostChecker as root; this will cause misleading or incorrect results from many of the checks.

4. Select which **checks** you want to run.

Run Tests without the Interface

You can also run checks without spawning the interface. Enter `--help` to get a list of arguments you can pass to the HostChecker.

5. As the checks are made, enter the requested **arguments**.
6. Read the output of the check.
The general format is that severity increases as you scroll down the output. First comes informational output, then warnings, then errors.

Internal Errors

If you see a message that starts with `Internal Error`, forward it to Delphix Support immediately. This represents a potential bug in the HostChecker, and not necessarily a problem with your environment.

7. Error or warning messages will explain any possible problems and how to address them. Resolve the issues that the HostChecker describes. Do not be surprised or undo your work if more errors appear the next time you run HostChecker, because the error you just fixed may have been masking other problems.
8. Repeat steps 3 - 7 until all the checks return no errors or warnings.

Adding Oracle Source and Target Environments

Follow the steps below to add **both** source and target environments for Oracle.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Plus** icon next to **Environments**.
5. In the **Add Environment** dialog, select **Unix/Linux**.
6. Select **Standalone Host** or **Oracle Cluster**, depending on the type of environment you are adding.
7. For standalone Oracle environments enter the **Host IP** address.
8. For Oracle RAC environments, enter the **Node Address** and **Cluster Home**.
9. Enter an optional **Name** for the environment.
10. Enter the **SSH** port.
The default value is **22**.
11. Enter a **Username** for the environment.
See [Requirements for Oracle Target Hosts and Databases](#) for more information on the required privileges for the environment user.
12. Select a Login Type:
 - Username and Password - enter the OS username and password, or
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key encryption for logging into your Unix-based environment:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

The public key needs to be added only once per user and per environment.

13. For **Password Login**, click **Verify Credentials** to test the username and password.
14. Enter a **Toolkit Path**.
The toolkit directory stores scripts used for Delphix Engine operations, and should have a persistent working directory rather than a temporary one. The toolkit directory will have a separate subdirectory for each database instance. The toolkit path must have 0770 permissions and at least 345MB of free space.
15. Click **Submit**.

Linking an Oracle Data Source

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.
4. In the **Add dSource** wizard, select the source database with the correct environment user-specified.
5. Enter your login credentials for the source database and click **Verify Credentials**.
If you are linking a mounted standby, Click **Next**. See the topics under [Linking Oracle Physical Standby Databases](#) for more information about how the Delphix Engine uses non-SYS login credentials.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed. For more information, visit [Data Management Settings for Oracle Data Sources](#).
8. Assign existing policies to the new dSource. New policies can be created and associated later.
9. Enter any scripts that should be run using the **Hooks** page.
10. Review the dSource Configuration and Data Management information, and then click **Submit**.

Once the action to add a dSource has been submitted, the Delphix Engine will initiate a DB_Link job to create the dSource. If the *Load Immediately* option was selected in the data management page a DB_Sync job will also be executed to ingest data from the source, otherwise, this first DB_Sync job will run as per the associated SnapSync policy.


When the jobs have successfully completed, the database icon will change to a dSource icon on the Environments > Databases screen, and the dSource will be added to the list of Datasets under its assigned group.

Provisioning an Oracle VDB

1. Login to the **Delphix Management** application.
2. Select **Manage > Datasets**.
3. In the **Datasets** panel on the left-hand side, click the **group** containing the dSource or VDB from which you want to provision.
4. Click the **TimeFlow** tab.
5. Select a **snapshot**.

 You can take a snapshot of the dSource from which to provision. To do so, click the **Camera** icon.

6. **Optional:** Select  to open LogSync timeline.



7. Select  to provision from a point of time within a snapshot. You can select by date or time.
8. Click  and the **Provision VDB** wizard will open:
 - a. For **Oracle Single Instance** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, and **Environment User** will auto-populate with information from the parent.
 - b. For **Oracle RAC** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, **Instance Number**, **Instance Name** and **Environment User** will auto-populate with information from the parent.

Editable Fields in the VDB Provision Wizard

The following fields are editable:

Installation Home (need to have an additional compatible target)

Database Unique Name

SID

Database Name

Mount Base

Instance Number (RAC Only)

Instance Name (RAC Only)

9. If you need to add a new target environment for the VDB, click the green **Plus** icon next to the **Filter Target** field, and follow the instructions in [Adding an Oracle Single Instance or RAC Environment](#).
10. Review the information for **Installation Home**, **Database Unique Name**, **SID**, and **Database Name**. Edit as necessary.
11. Review the **Mount Base** and **Environment User**. Edit as necessary.
The Environment User must have permission to write to the specified Mount Base, as described

in [Requirements for Oracle Target Hosts and Databases](#). You may also want to create a new writeable directory in the target environment with the correct permissions and use that as the Mount Base for the VDB.

12. Select **Provide Privileged Credentials** if you want to use login credentials on the target environment that are different from those associated with the **Environment User**.
13. Click **Advanced** to customize the VDB online log size and log groups, archivelog mode, local_listener parameter (TCP/IP protocol addresses), additional VDB configuration settings or file mappings, or custom environment variables. For more information, see [Customizing Oracle VDB Configuration Settings](#), [Customizing VDB File Mappings](#), and [Customizing Oracle VDB Environment Variables](#).



If you are provisioning to a target environment that is running a Linux OS, you will need to compare the `SGA_TARGET` configuration parameter with the shared memory size in `/dev/shm`. The shared memory configured on the target host should match the SGA memory target. You can check the Linux OS shared memory size with the command `df -k /dev/shm` and the `SGA_TARGET` configuration parameter by opening the **Advanced** settings and then finding the value for `SGA_TARGET` under **VDB Configuration**

Templates.

14. Click **Next**.
15. Select a **Target Group** for the VDB.
16. Enable Auto VDB Restart to allow the VDB to be automatically restarted when the target host reboot is detected by Delphix.
17. Click Next.
18. Select a **Snapshot Policy** for the VDB.
19. Click **Next**.
20. Enter any operations that should be run at Hooks during the provisioning process. Click **Next**.
21. Click **Submit**.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the **Status** tab, or by selecting **Manage/Dashboards** and viewing the **Job History** panel. Alternatively, you could see this in the **Actions Sidebar**. When provisioning is complete, the VDB will be included in the group you designated and listed in the **Datasets** panel. If you select the VDB in the **Datasets** panel and click the **Configuration** tab, you can view information about the database and its Data Management settings.

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.
2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test the sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision new VDBs from that masked VDB to quickly provide safe data to development and QA teams.

Script

- [createDelphixOSUser.sh](#)

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)
- [Provisioning and Managing Virtual Databases with Oracle](#)

Quick Start Guide for Oracle on Solaris SPARC


This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with Oracle database objects in the Delphix Engine. It does not cover advanced configuration options including Oracle RAC, Linking to Standby, or Best Practices for Performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix functionality. It assumes you will use the VMware Hypervisor.

Overview

In this guide, we will walk through deploying a Delphix Engine, configuring Oracle Source and Target environments on Solaris SPARC servers, creating a dSource, and provisioning a VDB.

The following diagram describes the Delphix topology for Oracle environments. It illustrates the recommended ports to be open from Delphix to remote services, to the Delphix Engine, and to the Source and Target Environments.

For this Quick Start Guide, you can ignore the following components: Validated Sync Host, Replicated Engine, and Delphix Masking Engine.



Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 similarly sized data disks for Delphix VM e.g. for 4 TB create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online.
This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
    example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
    the primary interface. Setting this value
                    to "true" will cause all other properties (address,
    hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
    "1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
    -- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24

```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
```

```
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

- Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, 192.168.1.2/24)

- Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

- Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the
browser at:
```

```
http://<address>
```


Type "exit" to disconnect, or any other commands to **continue** using the CLI.

- Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
- Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Once you setup the network access for Delphix Engine, enter Delphix Engine URL in your browser for server setup. The welcome screen below will appear for you to begin your Delphix Engine setup.



The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at <http://<Delphix Engine>/login/index.html#serverSetup>.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the delphix_admin are displayed for reference.

⚠ The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username '`delphix_admin`'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

System Time

Choose your option to setup system time in this section.

For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You've already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click Next to configure these for data storage.

Outbound Connectivity

Choose your options to configure outbound connectivity settings.

For a Quick Start, take the defaults. You can change this later.

Authentication

Choose your options to configure authentication services.

For a Quick Start, take the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.
4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

⚠ While your Delphix Engine will work without registration, we strongly recommend that you register each Delphix Engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at `http://<Delphix Engine>/login/index.html#serverSetup`.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Requirements for Oracle Hosts and Databases

Oracle Hosts and Databases

To begin using Oracle environments in Delphix, reference this article as a provided outline of the configuration requirements for the source and target hosts.

On each host with Oracle, there must be an operating system user configured to the required specifications for Delphix, as explained in the table below. This user (i.e. `delphix_os`) can easily be created by using the `createDelphixOSUser.sh` script (located at the bottom of the page).

These requirements apply to both source and target environments. However, target environments have additional requirements which are detailed in the 'Target Host Requirements' section below.

Source Host Requirement

Host Requirement	Explanation
Profile and privileges should be the same as the Oracle user (e.g. <code>oracle</code>) on the host.	For example, <code>delphix_os</code> should have the same environment variable settings (<code>\$PATH</code> , <code>\$ORACLE_HOME</code> , etc.), <code>umask</code> , and <code>ulimit</code> settings, as the user <code>oracle</code> .

<p>The Delphix software owner account (e.g. delphix_os) must have the same primary OS group as the Oracle software owner account (e.g. oracle).</p>	<p>Delphix recommends giving the delphix_os user the same primary OS group as the Oracle home owner. This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.</p> <p>Often, this is an OS group named oinstall. However, the oinstall group is not always necessary depending on your Oracle configuration. This user requires access to the libobk_proxy.so library in the toolkit for the child processes triggered by RMAN.</p>
<p>The delphix_os user must have the Oracle OSDBA group (typically dba) as a primary or secondary OS group.</p>	<p>The OSDBA group allows "OS authentication" when connecting to an Oracle database instance by specifying neither username or password (i.e. rman target /), thus eliminating the need to store or retrieve an SYSDBA password.</p> <p>Oracle 12c For Oracle 12c and later versions of Oracle, the delphix_os user can also use OSBACKUPDBA as its secondary group. This is typically the BACKUPDBA group on the host.</p>
<p>For secondary group requirements, the Delphix OS user must be a member of the SYSBACKUP OS group (12.1 or higher) or the SYSDBA OS group (11.2 and lower).</p>	<p>This ensures that the Delphix engine is able to take snapshots of source databases using RMAN.</p>
<p>There must be a directory on the source host where the Delphix Engine Toolkit can be installed, for example: /var/opt/delphix/toolkit</p>	<p>The delphix_os user and primary OS group (i.e. oinstall) must own the directory.</p> <p>The directory must have at least permissions -rwxr-x-- (0750)</p> <p>The delphix_os user must have -rwxr-x-- (0750) permissions on each directory in the path leading to the toolkit directory.</p> <p>At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.</p>
<p>The Delphix Engine must be able to make an SSH connection to the source host (typically port 22).</p>	
<p>Read access to \$ORACLE_HOME and all underlying files and directories.</p>	<p>Delphix needs to run locally available Oracle tools such as sqlplus and RMAN. Those executables, as well as various required libraries, reside inside the Oracle home</p>

Additional requirements for RAC Environments

If the source host is part of a RAC cluster, Delphix will attempt to use all nodes and crsctl for its operations.

RAC Environment Requirement	Explanation
The delphix_os user and Delphix Toolkit configuration must be the same on each node in the RAC cluster	<ul style="list-style-type: none"> delphix_os must have a execute permission on crsctl and srvctl on each node in the cluster. access to olsnodes is needed. The Delphix Toolkit must be installed in the same directory path on each of the nodes in the source cluster. All data files, archive logs, and database control file must be located on storage shared by all of the cluster nodes. Each node in the cluster must be able to access archive logs and the database control file from all other nodes. At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.

Masked Provisioning

Masked Provisioning is supported on Oracle RAC only when used with "script-based masking".

Auto-Discovery Requirements

The preferred way to find source databases is to allow Delphix to automatically discover your Oracle Homes and Databases by examining the inventory and oratab files and the listener setup. Successful autodiscovery requires read access to these and related files.

In **most** environments, delphix_os group membership is sufficient to perform auto-discovery.

However, if you have overridden Oracle's group permission structure, you may need to modify privileges to allow auto-discovery.

Auto-Discovery Requirements	Explanation
<p>The ORATAB file must exist (typically in /etc/oratab or /var/opt/oracle/oratab) and be readable by delphix_os.</p> <p>Read access to:</p> <ul style="list-style-type: none"> /etc/orainst.loc or /var/opt/oracle/orainst.loc. the Oracle inventory file (inventory.xml) 	<p>The ORATAB is used to determine the location of the Oracle installations on your host.</p> <p>Note: The Oracle inventory file is identified by the contents of oraInst.loc (for example, \$INVENTORY_HOME/ContentsXML/inventory.xml).</p>

<p>Permission to run pargs on Solaris hosts and ps on AIX, HP-UX, and Linux hosts, as super-user.</p> <p>This permission is usually granted via sudo authorization of the commands.</p>	<p>Unless you have used a custom TNS_ADMIN setting, elevated access to ps (pargs on Solaris) is not required</p> <p>See the topic Sudo Privilege Requirements for Oracle Environments for further explanation of this requirement, and Sudo File Configuration Examples for Oracle Environments for examples of file configurations</p>
---	---

Source Database Requirements

For each source database, there are specific configurations required for Delphix to ingest data.

Source Database Requirement	Explanation
For each Oracle Home, the delphix_os user should have execute permission for the programs in \$ORACLE_HOME/bin.	<p>If symlinks are configured, Delphix must be configured with the same \$ORACLE_HOME path as was used when starting the instance.</p> <p>Ensure the PermitUserEnvironment configuration parameter = "yes" in the sshd_config file</p> <p>The \$ORACLE_HOME/bin/oracle executable must have the SETUID and SETGID flags set. Permissions on the oracle binary must be at least -rwsr-s-x (06751).</p>
Source databases must be in ARCHIVELOG mode to ensure that redo logs are archived.	Archive logs are required to make SnapSyncs consistent and provisionable.
Enable Block Change Tracking (BCT). (Highly Recommended).	Enabling BCT will improve SnapSync operation time. Without BCT, incremental SnapSyncs must scan the entire database.
Enable FORCE LOGGING. (Highly Recommended)	<p>If you do not enable FORCE LOGGING and NOLOGGING operations take place, you will get a Fault from the Delphix Engine.</p> <p>If you must use NOLOGGING to meet specific performance criteria, take a new snapshot of the source database after doing the NOLOGGING operations to bring the dSource up-to-date before provisioning VDBs.</p>
If the online redo log files are located on RAW or ASM devices, then LogSync can only operate in Archive Only mode.	See the topics Advanced Data Management Settings for Oracle dSources and Linking Oracle Physical Standby Databases for more information.

Operating System Specific Requirements

Solaris

On a Solaris host, gtar must be installed. Delphix uses gtar to handle long file names when extracting the toolkit files into the toolkit directory on a Solaris host. The gtar binary should be installed in one of the following directories:

- / bin:/usr

- / bin:/sbin:/usr
- / sbin:/usr/contrib
- / bin:/usr/sfw
- / bin:/opt/sfw
- / bin:/opt/csw/bin

Additional Target Host Requirements

This topic describes the user privileges and environment requirements that are required for Oracle target hosts and databases collectively referred to as target environments.

These are in addition to the ‘Oracle Hosts and Database Requirements’ called out above.

Target Host Requirement	Explanation
The Delphix OS user must be a member of the SYSDBA OS group.	This ensures that the Delphix engine is able to create new VDBs on target hosts.
There must be a directory (e.g. "/mnt/provision/") that will be used as a container for the NFS mount points that are created when provisioning a VDB.	<p>The delphix_os user and primary OS group (i.e. oinstall or dba) must own the directory.</p> <p>The directory must have at least permissions -rwxrwx--- (0770).</p> <p>There must be no symbolic links in the path of this directory, because NFS can mount into a directory with symlinks in its path, but cannot unmount.</p>
Permission to run: <ul style="list-style-type: none"> • mount, umount, mkdir, rmdir as super-user. • pargs on Solaris hosts • ps on Linux, AIX, and HP-UX as super-user. • nfso on AIX hosts as super-user 	The following permissions are usually granted via sudo authorization of the commands. See Sudo Privilege Requirements for Oracle Environments for further explanation of the commands, and Sudo File Configuration Examples for Oracle Environments for examples of the /etc/sudoers file on different operating systems.
Optional write permission to the \$ORACLE_BASE_CONFIG/dbs directory	<p>If write permission is granted, an instance init file will be written to this directory during provisioning that will simplify manual instance startup.</p> <p>If write permission is not granted, manual instance startup must specify the instance init file. See Manually Starting a VDB.</p>
An Oracle listener process running on the target host.	The listener's version should be equal to or greater than the highest Oracle version that will be used to provision a VDB.
Required packages on target hosts: <ol style="list-style-type: none"> 1. portmapper / rpcbind 2. status daemon (rpc.statd) 3. NFS lock manager (rpc.lockd/lockmgr) 	As the Delphix Engine uses NFSv3 for mounting target host filesystems, the prerequisite packages to support NFSv3 client communication are required for normal operation, and the required services to support NFS client communications (including file locking) must be running. NFSv3 is enabled by default and to enable NFSv4, see NFSv4 Configuration .

If using ASM, the `delphix_os` user must have the Oracle ASM ownership groups as secondary OS groups (typically `asmadmin` and `asmdba`)

This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.

Deploy Hostchecker to Validate Delphix Requirements

Delphix has developed a hostchecker script that contains standardized checks for source and target hosts - these checks generally fall into three buckets

- OS and Host permissions/access
- Network Port Checks
- DB-specific functionality

OS and Host permissions/access and network port checks can (and should) be performed prior to Delphix installation to ensure a smooth deployment.

Each DB should have a specifically associated hostchecker - you can find detailed documentation on the DB-specific hostchecker page.

1. Download the appropriate **HostChecker tarball** for your engine from <https://download.delphix.com/>. Tarballs follow the naming convention "hostchecker_<OS>_<processor>.tar". For example, if you are validating a linux x86 host you should download the `hostchecker_linux_x86.tar` tarball.
2. Create a working directory and extract the **HostChecker files** from the **HostChecker tarball**.

```
mkdir dlp-host-checker
cd dlp-host-checker/
tar -xf hostchecker_linux_x86.tar
```

3. Run the `sh` script contained within:

```
sh hostchecker.sh
```

This will extract the JDK included in the tarball (if necessary) and invoke the HostChecker.

```
ora10205@bbdhcp:/home/ora10205/hostchecker-> sh hostchecker.sh
Extracting the JDK from the tarball jdk-6u45-linux-i586.tar.gz.
```

Don't Run as Root

Do not run the HostChecker as root; this will cause misleading or incorrect results from many of the checks.

4. Select which **checks** you want to run.

Run Tests without the Interface

You can also run checks without spawning the interface. Enter `--help` to get a list of arguments you can pass to the HostChecker.

5. As the checks are made, enter the requested **arguments**.

6. Read the output of the check.
The general format is that severity increases as you scroll down the output. First comes informational output, then warnings, then errors.

Internal Errors

If you see a message that starts with `Internal Error`, forward it to Delphix Support immediately. This represents a potential bug in the HostChecker, and not necessarily a problem with your environment.

7. Error or warning messages will explain any possible problems and how to address them. Resolve the issues that the HostChecker describes. Do not be surprised or undo your work if more errors appear the next time you run HostChecker, because the error you just fixed may have been masking other problems.
8. Repeat steps 3 - 7 until all the checks return no errors or warnings.

Adding Oracle Source and Target Environments

Follow the steps below to add **both** source and target environments for Oracle.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Plus** icon next to **Environments**.
5. In the **Add Environment** dialog, select **Unix/Linux**.
6. Select **Standalone Host** or **Oracle Cluster**, depending on the type of environment you are adding.
7. For standalone Oracle environments enter the **Host IP** address.
8. For Oracle RAC environments, enter the **Node Address** and **Cluster Home**.
9. Enter an optional **Name** for the environment.
10. Enter the **SSH** port.
The default value is **22**.
11. Enter a **Username** for the environment.
See [Requirements for Oracle Target Hosts and Databases](#) for more information on the required privileges for the environment user.
12. Select a Login Type:
 - Username and Password - enter the OS username and password, or
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key encryption for logging into your Unix-based environment:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

The public key needs to be added only once per user and per environment.

13. For **Password Login**, click **Verify Credentials** to test the username and password.

14. Enter a **Toolkit Path**.

The toolkit directory stores scripts used for Delphix Engine operations, and should have a persistent working directory rather than a temporary one. The toolkit directory will have a separate subdirectory for each database instance. The toolkit path must have 0770 permissions and at least 345MB of free space.

15. Click **Submit**.

Linking an Oracle Data Source

1. Login to the **Delphix Management** application.2. Navigate to **Manage > Datasets**.3. Click the plus icon and select **Add dSource**.4. In the **Add dSource** wizard, select the source database with the correct environment user-specified.5. Enter your login credentials for the source database and click **Verify Credentials**.

If you are linking a mounted standby, Click **Next**. See the topics under [Linking Oracle Physical Standby Databases](#) for more information about how the Delphix Engine uses non-SYS login credentials.

6. Enter a name and select a group for your dSource.

Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.

7. Select the **Data Management** settings needed. For more information, visit [Data Management Settings for Oracle Data Sources](#).

8. Assign existing policies to the new dSource. New policies can be created and associated later.

9. Enter any scripts that should be run using the **Hooks** page.10. Review the dSource Configuration and Data Management information, and then click **Submit**.

Once the action to add a dSource has been submitted, the Delphix Engine will initiate a DB_Link job to create the dSource. If the *Load Immediately* option was selected in the data management page a DB_Sync job will also be executed to ingest data from the source, otherwise, this first DB_Sync job will run as per the associated SnapSync policy.


When the jobs have successfully completed, the database icon will change to a dSource icon on the Environments > Databases screen, and the dSource will be added to the list of Datasets under its assigned group.

Provisioning an Oracle VDB

1. Login to the **Delphix Management** application.2. Select **Manage > Datasets**.3. In the **Datasets** panel on the left-hand side, click the **group** containing the dSource or VDB from which you want to provision.4. Click the **TimeFlow** tab.5. Select a **snapshot**.

You can take a snapshot of the dSource from which to provision. To do so, click the **Camera** icon.

6. **Optional:** Select  to open LogSync timeline.



7. Select  to provision from a point of time within a snapshot. You can select by date or time.
8. Click  and the **Provision VDB** wizard will open:
 - a. For **Oracle Single Instance** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, and **Environment User** will auto-populate with information from the parent.
 - b. For **Oracle RAC** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, **Instance Number**, **Instance Name** and **Environment User** will auto-populate with information from the parent.

Editable Fields in the VDB Provision Wizard

The following fields are editable:

Installation Home (need to have an additional compatible target)

Database Unique Name

SID

Database Name

Mount Base

Instance Number (RAC Only)

Instance Name (RAC Only)

9. If you need to add a new target environment for the VDB, click the green **Plus** icon next to the **Filter Target** field, and follow the instructions in [Adding an Oracle Single Instance or RAC Environment](#).
10. Review the information for **Installation Home**, **Database Unique Name**, **SID**, and **Database Name**. Edit as necessary.
11. Review the **Mount Base** and **Environment User**. Edit as necessary.
The Environment User must have permission to write to the specified Mount Base, as described

in [Requirements for Oracle Target Hosts and Databases](#). You may also want to create a new writeable directory in the target environment with the correct permissions and use that as the Mount Base for the VDB.

12. Select **Provide Privileged Credentials** if you want to use login credentials on the target environment that are different from those associated with the **Environment User**.
13. Click **Advanced** to customize the VDB online log size and log groups, archivelog mode, local_listener parameter (TCP/IP protocol addresses), additional VDB configuration settings or file mappings, or custom environment variables. For more information, see [Customizing Oracle VDB Configuration Settings](#), [Customizing VDB File Mappings](#), and [Customizing Oracle VDB Environment Variables](#).



If you are provisioning to a target environment that is running a Linux OS, you will need to compare the `SGA_TARGET` configuration parameter with the shared memory size in `/dev/shm`. The shared memory configured on the target host should match the SGA memory target. You can check the Linux OS shared memory size with the command `df -k /dev/shm` and the `SGA_TARGET` configuration parameter by opening the **Advanced** settings and then finding the value for `SGA_TARGET` under **VDB Configuration**

Templates.

14. Click **Next**.
15. Select a **Target Group** for the VDB.
16. Enable Auto VDB Restart to allow the VDB to be automatically restarted when the target host reboot is detected by Delphix.
17. Click Next.
18. Select a **Snapshot Policy** for the VDB.
19. Click **Next**.
20. Enter any operations that should be run at Hooks during the provisioning process. Click **Next**.
21. Click **Submit**.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the **Status** tab, or by selecting **Manage/Dashboards** and viewing the **Job History** panel. Alternatively, you could see this in the **Actions Sidebar**. When provisioning is complete, the VDB will be included in the group you designated and listed in the **Datasets** panel. If you select the VDB in the **Datasets** panel and click the **Configuration** tab, you can view information about the database and its Data Management settings.

Next Steps

Congratulations! You've provisioned your first virtual database!

You should attempt some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the Target listener.

We suggest the following next steps:

- dropping a table and using the VDB Rewind feature to test recovery of your VDB
- Take a snapshot of your dSource and Refresh your VDB to quickly get fresh production data
- Provision a new VDB from your VDB to test sharing data quickly with other teams
- Mask your new VDB to protect sensitive data. Provision new VDBs from that Masked VDB to quickly provide safe data to development and QA teams.

Script

- [createDelphixOSUser.sh](#)

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)
- [Provisioning and Managing Virtual Databases with Oracle](#)

Oracle Support and Requirements

This section covers the following topics:

- [Oracle Support Matrix](#)
- [Requirements for Oracle Hosts and Databases](#)
- [Network and Connectivity Requirements for Oracle Environments](#)
- [Sudo Privilege Requirements for Oracle Environments](#)
- [Sudo File Configuration Examples for Oracle Environments](#)

Oracle Support Matrix

Disclaimers and Support Notes

Delphix Support Policies specifically list **major** and **minor** release coverage. If a minor release is listed as covered, all patch releases under that minor release are certified. Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL), 64-bit OS support only.

The Oracle version in the support matrix applies to both Oracle DB and Oracle Grid. Delphix supports the same set of features and functionality for supported non-multitenant database versions. Currently, Delphix does not support the Oracle 12c feature of `THREADED_EXECUTION` being set to `TRUE`, because this disables OS authentication.

Summary of Delphix features that are unsupported for the Oracle12c and higher Multi-tenant configuration:

- Existing Virtual Container Database (vCDB) as target for provisioning an additional vPDB
- Customize VDB settings/initialization parameters. Includes the following:
 - Customize `init.ora` database parameter during provisioning
 - Config templates
 - Online redo log size
 - Number of RAC VDB instances
 - Online redo log groups
 - Archive log mode
 - Setting new DBID
 - Customize local listeners
- Virtual to Physical (V2P) Support
- Resumable initial SnapSync
- Validated Sync
- Source continuity for dSource upgraded from Oracle 12c non-multitenant to multitenant database
- Cross-platform provisioning (XPP) to the virtual database
- Oracle LiveSources

Delphix supports systems that are part of an Exadata or Exadata Cloud-at-Customer (ExaCC) cluster, provided that the operating system and DBMS version are in the supported list. Oracle Exadata, Exadata Cloud, and Exadata Cloud-at-Customer (ExaCC) offer several features that are reliant on the underlying Exadata storage. Delphix supports creating dSources from Exadata source clusters and provisioning VDBs and vPDBs to Exadata target clusters.

The following Exadata-only features will not be available in child VDBs running in Delphix storage:

- Smart Scan - allows SQL processing to be offloaded to Exadata storage cells. Queries that leverage Smart Scan will execute successfully in VDBs but query performance may be impacted.
- Smart Flash Cache - allows frequently used data blocks to be stored in the flash cache of Exadata storage cells to allow for faster access. Queries against objects that are normally cached in Smart Flash Cache will execute successfully in VDBs but query performance may be impacted. As the virtualization engine uses memory to cache frequently used data blocks, allocating extra memory to the engine may help to mitigate this performance impact.
- Hybrid Columnar Compression (HCC) - compresses data and stores it in a columnar format, allowing Exadata storage cells to deliver faster query performance while also saving on storage. Columns compressed using HCC cannot be used by queries executed in VDBs. Affected columns can be decompressed though care should be taken to evaluate the impact of doing so on VDB refresh times and Delphix storage usage. In Delphix VDBs, queries against partitions or tables using HCC will fail with the following error *"ORA-64307: Exadata Hybrid Columnar Compression is not supported for tablespaces on this storage type"*.

Color	Supported?
	Yes
	No
	Not Applicable

Oracle Database Editions

Delphix supports the following Oracle database editions:

- Enterprise
- Standard

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please refer to [KBA 5622](#) for the required steps.

Supported OS Version	Supported DBMS Version						
	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 5.5							
RHEL 5.6							
RHEL 5.7							
RHEL 5.8							
RHEL 5.9							
RHEL 5.10							
RHEL 5.11							

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 6.0	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.1	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.2	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.3	N/A	✓	✓	N/A	N/A	N/A	N/A
RHEL 6.4	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.5	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.6	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.7	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.8	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.9	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 6.10	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.0	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.1	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.2	N/A	✓	✓	✓	✓	N/A	N/A
RHEL 7.3	N/A	✓	✓	✓	✓	N/A	N/A

Supported OS Version	Supported DBMS Version						
	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
RHEL 7.4	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.5	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.6	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.7	N/A	✓	✓	✓	✓	✓	N/A
RHEL 7.8	N/A	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	✓ 6.0.2+	N/A
RHEL 7.9	N/A	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.4+	✓ 6.0.5+	✓ 6.0.4+	N/A
RHEL 8.0	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.1	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.2	N/A	N/A	N/A	N/A	N/A	✓ 6.0.3+	N/A
RHEL 8.3	N/A	N/A	N/A	N/A	N/A	✓ 6.0.7+	✓ 6.0.11+
RHEL 8.4	N/A	N/A	N/A	N/A	N/A	✓ 6.0.10+	✓ 6.0.11+

SUSE Linux Enterprise Server (SLES)



Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 11	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP1	✓	✓	N/A	N/A	N/A	N/A	N/A
SLES 11 SP2	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP3	✓	✓	✓	N/A	N/A	N/A	N/A
SLES 11 SP4	✗	✓	✓	N/A	N/A	N/A	N/A
SLES 12	N/A	✓	✓	N/A	N/A	N/A	N/A
SLES 12 SP1	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP2	N/A	✓	✓	✓	✓	N/A	N/A
SLES 12 SP3	N/A	✓	✓	✓	✓	✓	N/A
SLES 12 SP4	N/A	✓	✓	✓	✓	✓	N/A
SLES 12 SP5	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+	✗	✓ 6.0.4+	N/A
SLES 15	N/A	✗	✗	✓	✓	✓	N/A
SLES 15 SP1	N/A	✗	✗	✗	✗	✓	✓ 6.0.11 +
SLES 15 SP2	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11 +

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
SLES 15 SP3	N/A	✗	✗	✗	✗	✓ 6.0.11+	✓ 6.0.11+

Solaris SPARC

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	✗	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i ¹ - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 SPARC kernel patch 144500, or a newer kernel patch.

Solaris x86

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
Solaris 10 U9 ¹	✓	✓	N/A	N/A	N/A	N/A	N/A
Solaris 10 U10 ¹	✓	✓	✓	N/A	N/A	N/A	N/A
Solaris 10 U11	✓	✓	✓	N/A	✓	N/A	N/A
Solaris 11	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U1	N/A	✓	✓	N/A	N/A	N/A	N/A
Solaris 11 U2	N/A	✓	✓	✓	✓	N/A	N/A
Solaris 11 U3	N/A	✗	✓	✓	✓	✓	N/A
Solaris 11 U4	N/A	✗	✗	✓	✓	✓	N/A

i ¹ - Solaris 10 U9 and U10 require libc.so.1 version 1.22.7 or newer. This version of libc.so.1 can be found in Solaris 10 X86 kernel patch 144501, or a newer kernel patch.

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
HP-UX 11.31	✓	✓	✓	✓	✗	✗	N/A

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version						
Supported OS Version	11gR1	11gR2	12cR1	12cR2	18c	19c	21c
AIX 6.1	✓	✓	✓	✗	✗	N/A	N/A
AIX 7.1	N/A	✓	✓	✓	✓	✓	N/A
AIX 7.2	N/A	✓	✓	✓	✓	✓	N/A

Related Topics

- [Oracle Support and Requirements](#)

Requirements for Oracle Hosts and Databases

Oracle Hosts and Databases

To begin using Oracle environments in Delphix, reference this article as a provided outline of the configuration requirements for the source and target hosts.

On each host with Oracle, there must be an operating system user configured to the required specifications for Delphix, as explained in the table below. This user (i.e. `delphix_os`) can easily be created by using the `createDelphixOSUser.sh` script (located at the bottom of the page).

These requirements apply to both source and target environments. However, target environments have additional requirements which are detailed in the 'Target Host Requirements' section below.

Source Host Requirement

Host Requirement	Explanation
Profile and privileges should be the same as the Oracle user (e.g. <code>oracle</code>) on the host.	For example, <code>delphix_os</code> should have the same environment variable settings (<code>\$PATH</code> , <code>\$ORACLE_HOME</code> , etc.), <code>umask</code> , and <code>ulimit</code> settings, as the user <code>oracle</code> .
The Delphix software owner account (e.g. <code>delphix_os</code>) must have the same primary OS group as the Oracle software owner account (e.g. <code>oracle</code>).	<p>Delphix recommends giving the <code>delphix_os</code> user the same primary OS group as the Oracle home owner. This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.</p> <p>Often, this is an OS group named <code>oinstall</code>. However, the <code>oinstall</code> group is not always necessary depending on your Oracle configuration. This user requires access to the libobk_proxy.so library in the toolkit for the child processes triggered by RMAN.</p>
The <code>delphix_os</code> user must have the Oracle OSDBA group (typically <code>dba</code>) as a primary or secondary OS group.	<p>The OSDBA group allows "OS authentication" when connecting to an Oracle database instance by specifying neither username or password (i.e. <code>rman target /</code>), thus eliminating the need to store or retrieve an SYSDBA password.</p> <p>Oracle 12c For Oracle 12c and later versions of Oracle, the <code>delphix_os</code> user can also use OSBACKUPDBA as its secondary group. This is typically the BACKUPDBA group on the host.</p>
For secondary group requirements, the Delphix OS user must be a member of the SYSBACKUP OS group (12.1 or higher) or the SYSDBA OS group (11.2 and lower).	This ensures that the Delphix engine is able to take snapshots of source databases using RMAN.

There must be a directory on the source host where the Delphix Engine Toolkit can be installed, for example: /var/opt/delphix/toolkit	<p>The delphix_os user and primary OS group (i.e. oinstall) must own the directory.</p> <p>The directory must have at least permissions -rwxr-x-- (0750)</p> <p>The delphix_os user must have -rwxr-x-- (0750) permissions on each directory in the path leading to the toolkit directory.</p> <p>At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.</p>
The Delphix Engine must be able to make an SSH connection to the source host (typically port 22).	
Read access to \$ORACLE_HOME and all underlying files and directories.	Delphix needs to run locally available Oracle tools such as sqlplus and RMAN. Those executables, as well as various required libraries, reside inside the Oracle home

Additional requirements for RAC Environments

If the source host is part of a RAC cluster, Delphix will attempt to use all nodes and crsctl for its operations.

RAC Environment Requirement	Explanation
The delphix_os user and Delphix Toolkit configuration must be the same on each node in the RAC cluster	<ul style="list-style-type: none"> delphix_os must have a execute permission on crsctl and srvctl on each node in the cluster. access to olsnodes is needed. The Delphix Toolkit must be installed in the same directory path on each of the nodes in the source cluster. All data files, archive logs, and database control file must be located on storage shared by all of the cluster nodes. Each node in the cluster must be able to access archive logs and the database control file from all other nodes. At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.

Masked Provisioning

Masked Provisioning is supported on Oracle RAC only when used with "script-based masking".

Auto-Discovery Requirements

The preferred way to find source databases is to allow Delphix to automatically discover your Oracle Homes and Databases by examining the inventory and oratab files and the listener setup. Successful autodiscovery requires read access to these and related files.

In **most** environments, `delphix_os` group membership is sufficient to perform auto-discovery.

However, if you have overridden Oracle's group permission structure, you may need to modify privileges to allow auto-discovery.

Auto-Discovery Requirements	Explanation
<p>The ORATAB file must exist (typically in <code>/etc/oratab</code> or <code>/var/opt/oracle/oratab</code>) and be readable by <code>delphix_os</code>.</p> <p>Read access to:</p> <ul style="list-style-type: none"> <code>/etc/orainst.loc</code> or <code>/var/opt/oracle/orainst.loc</code>. the Oracle inventory file (<code>inventory.xml</code>) 	<p>The ORATAB is used to determine the location of the Oracle installations on your host.</p> <p>Note: The Oracle inventory file is identified by the contents of <code>orainst.loc</code> (for example, <code>\$INVENTORY_HOME/ContentsXML/inventory.xml</code>).</p>
<p>Permission to run <code>pargs</code> on Solaris hosts and <code>ps</code> on AIX, HP-UX, and Linux hosts, as <code>super-user</code>.</p> <p>This permission is usually granted via <code>sudo</code> authorization of the commands.</p>	<p>Unless you have used a custom <code>TNS_ADMIN</code> setting, elevated access to <code>ps</code> (<code>pargs</code> on Solaris) is not required</p> <p>See the topic Sudo Privilege Requirements for Oracle Environments for further explanation of this requirement, and Sudo File Configuration Examples for Oracle Environments for examples of file configurations</p>

Source Database Requirements

For each source database, there are specific configurations required for Delphix to ingest data.

Source Database Requirement	Explanation
<p>For each Oracle Home, the <code>delphix_os</code> user should have execute permission for the programs in <code>\$ORACLE_HOME/bin</code>.</p>	<p>If symlinks are configured, Delphix must be configured with the same <code>\$ORACLE_HOME</code> path as was used when starting the instance.</p> <p>Ensure the <code>PermitUserEnvironment</code> configuration parameter = "yes" in the <code>sshd_config</code> file</p> <p>The <code>\$ORACLE_HOME/bin/oracle</code> executable must have the <code>SETUID</code> and <code>SETGID</code> flags set. Permissions on the oracle binary must be at least <code>-rwsr-s-x</code> (06751).</p>
<p>Source databases must be in <code>ARCHIVELOG</code> mode to ensure that redo logs are archived.</p>	<p>Archive logs are required to make SnapSyncs consistent and provisionable.</p>
<p>Enable Block Change Tracking (BCT). (Highly Recommended).</p>	<p>Enabling BCT will improve SnapSync operation time. Without BCT, incremental SnapSyncs must scan the entire database.</p>

Enable FORCE LOGGING. (Highly Recommended)	<p>If you do not enable FORCE LOGGING and NOLOGGING operations take place, you will get a Fault from the Delphix Engine.</p> <p>If you must use NOLOGGING to meet specific performance criteria, take a new snapshot of the source database after doing the NOLOGGING operations to bring the dSource up-to-date before provisioning VDBs.</p>
If the online redo log files are located on RAW or ASM devices, then LogSync can only operate in Archive Only mode.	See the topics Advanced Data Management Settings for Oracle dSources and Linking Oracle Physical Standby Databases for more information.

Operating System Specific Requirements

Solaris

On a Solaris host, gtar must be installed. Delphix uses gtar to handle long file names when extracting the toolkit files into the toolkit directory on a Solaris host. The gtar binary should be installed in one of the following directories:

- / bin:/usr
- / bin:/sbin:/usr
- / sbin:/usr/contrib
- / bin:/usr/sfw
- / bin:/opt/sfw
- / bin:/opt/csw/bin

Additional Target Host Requirements

This topic describes the user privileges and environment requirements that are required for Oracle target hosts and databases collectively referred to as target environments.

These are in addition to the ‘Oracle Hosts and Database Requirements’ called out above.

Target Host Requirement	Explanation
The Delphix OS user must be a member of the SYSDBA OS group.	This ensures that the Delphix engine is able to create new VDBs on target hosts.
There must be a directory (e.g. "/mnt/provision/") that will be used as a container for the NFS mount points that are created when provisioning a VDB.	<p>The delphix_os user and primary OS group (i.e. oinstall or dba) must own the directory.</p> <p>The directory must have at least permissions -rwxrwx--- (0770).</p> <p>There must be no symbolic links in the path of this directory, because NFS can mount into a directory with symlinks in its path, but cannot unmount.</p>

<p>Permission to run:</p> <ul style="list-style-type: none"> • mount, umount, mkdir, rmdir as super-user. • pargs on Solaris hosts • ps on Linux, AIX, and HP-UX as super-user. • nfso on AIX hosts as super-user 	<p>The following permissions are usually granted via sudo authorization of the commands. See Sudo Privilege Requirements for Oracle Environments for further explanation of the commands, and Sudo File Configuration Examples for Oracle Environments for examples of the /etc/sudoers file on different operating systems.</p>
<p>Optional write permission to the \$ORACLE_BASE_CONFIG/dbs directory</p>	<p>If write permission is granted, an instance init file will be written to this directory during provisioning that will simplify manual instance startup. If write permission is not granted, manual instance startup must specify the instance init file. See Manually Starting a VDB.</p>
<p>An Oracle listener process running on the target host.</p>	<p>The listener's version should be equal to or greater than the highest Oracle version that will be used to provision a VDB.</p>
<p>Required packages on target hosts:</p> <ol style="list-style-type: none"> 1. portmapper / rpcbind 2. status daemon (rpc.statd) 3. NFS lock manager (rpc.lockd/lockmgr) 	<p>As the Delphix Engine uses NFSv3 for mounting target host filesystems, the prerequisite packages to support NFSv3 client communication are required for normal operation, and the required services to support NFS client communications (including file locking) must be running. NFSv3 is enabled by default and to enable NFSv4, see NFSv4 Configuration.</p>
<p>If using ASM, the delphix_os user must have the Oracle ASM ownership groups as secondary OS groups (typically asmadmin and asmdba)</p>	<p>This ensures the Delphix engine can fully and automatically discover Oracle homes, databases, and listeners.</p>

Scripts

- [createDelphixOSUser.sh](#)
- [createDelphixDBUser.sh](#)

Related Topics

- [Network and Connectivity Requirements for Oracle Environments](#)
- [Sudo Privilege Requirements for Oracle Environments](#)
- [Sudo File Configuration Examples for Oracle Environments](#)

Network and Connectivity Requirements for Oracle Environments

General Port Allocation

The Delphix Engine makes use of the following network ports regardless of the type of database platform:

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine
TCP	443	HTTPS connections to the Delphix Management Application

Protocol	Port Number	Use
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

SSHD Configuration

Both source and target Unix environments are required to have `sshd` running and configured such that the Delphix Engine can connect over `ssh`.

The Delphix platform expects to maintain long-running, highly performant `ssh` connections with remote Unix environments. The following `sshd` configuration entries can interfere with these `ssh` connections and are therefore disallowed:

Disallowed <code>sshd</code> Configuration Entries
<code>ClientAliveInterval</code>
<code>ClientAliveCountMax</code>


Network and Connectivity Requirements for Oracle

- IP connections must exist between the Delphix Engine and the source and target environments.
- For source environments, Delphix Engine uses an **SSH** connection to each source host, an **HTTP** connection from each source environment to Delphix Engine, and a DSP connection to the Delphix Engine. The Delphix Engine uses **SQL*Net** connections to the DBMS on the source environment.
- For target environments, Delphix uses an **SSH** connection to each target environment and an **NFS** connection to Delphix Engine. Delphix Engine uses **SQL*Net** connections to the virtual databases on the target environment.

⚠ scp Availability
 The **scp** program must be available in the environment in order to add an environment.

Port Allocation for Oracle Environments

The following diagram describes the port allocations for Oracle environments. It illustrates the ports that we recommend to be open from Delphix to remote services, to the Delphix Engine, and to the Target Environments.



The Delphix Engine makes use of the following network ports for Oracle dSources and VDBs:

Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to source and target environments
TCP	xxx	Connections to the Oracle SQL*Net Listener on the source and target environments (typically port 1521)

Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP/UDP	111	Remote Procedure Call (RPC) port mapper used for NFSv3 mounts Note: RPC calls in NFSv3 use additional fixed ports for supporting services (lockd, mountd and statd) seen below.
TCP	1110	NFSv3 Server daemon status and NFS server daemon keep-alive (client info)
TCP/UDP	2049	NFS Server daemon from VDB to the Delphix Engine (used for both NFSv3 and NFSv4)
TCP	8341	Sending data from source to the Delphix Engine (for LogSync)
TCP	8415	SnapSync control and data from source to the Delphix Engine V2P control and data from the target environment to the Delphix Engine.
TCP	54043	NFSv3 mount daemon
TCP	54044	NFSv3 stat daemon (lock state notification service)
TCP	54045	NFSv3 lock daemon/manager
UDP	33434 - 33464	Traceroute from source and target database servers to the Delphix Engine (optional)

AppData Port Requirements

The use of AppData requires the following ports/protocols.

Two important notes about these specifications:

1. The next release of the Delphix Engine will significantly augment the port/protocol utilization of AppData. The upcoming-only requirements have been marked with a *.
2. AppData V2P uses RSYNC to export to the target. RSYNC between the target and Delphix Engine is not required for general virtualization usage. The V2P-only requirements have been marked with a ^.

From Source to Delphix Engine	From Delphix Engine to Source	From Target to Delphix Engine	From Delphix Engine to Target
RSYNC (TCP Port 873)	RSYNC (TCP Port 873)	DSP (Default TCP Port 8415)	DSP (Default TCP Port 8415)
DSP (Default TCP Port 8415)	SSH (TCP Port 22)	NFS	SSH (TCP Port 22)
*NFS	DSP (Default TCP Port 8415)	^RSYNC (TCP Port 873)	^RSYNC (TCP Port 873)

Related Topics

- [Configuring Replication](#)
- [Requirements for Oracle Hosts and Databases](#)

Sudo Privilege Requirements for Oracle Environments

This topic describes the rationale behind specific `sudo` privilege requirements for virtualizing Oracle Databases.

Privilege	Sources	Targets	Rationale
<code>ps pargs</code>	Optional, Strongly Recommended	Optional, Strongly Recommended	<p>Delphix auto-discovery uses the <code>TNS_ADMIN</code> environment variable of Oracle Listener processes with non-standard configurations to derive their connection parameters. A different user (oracle), rather than delphix_os user normally owns the oracle listener. The Delphix Engine needs <code>sudo</code> access to <code>pargs</code> on the Solaris OS or <code>ps</code> on other OSes to determine the environment variables of those Listener processes.</p> <p>This privilege is required for Auto-Discovery with non-default <code>TNS_ADMIN</code> locations. It is optional when using a standard <code>TNS_ADMIN</code> location, or if you choose to manually configure Oracle Homes, databases and listeners.</p>
<code>mount/umount</code>	Not Required	Required	<p>The Delphix Engine dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for superuser.</p>
<code>nfso</code> (AIX only)	Not Required	Required	<p>The Delphix Engine monitors NFS read and write sizes on an AIX target host. It uses the <code>nfso</code> command to query the sizes in order to optimize NFS performance for VDBs running on the target host. Only a superuser can issue the <code>nfso</code> command.</p>

⚠ It is required to specify the NOPASSWD qualifier within the "sudo" configuration file, as shown here: [Sudo File Configuration Examples for Oracle Environments](#). This ensures that the "sudo" command does not demand the entry of a password, even for the "display permissions" (i.e. "sudo -l") command. Delphix issues "sudo -l" in some scripts to detect if the operating system user has the correct sudo privileges. If it is unable to execute this command, some actions may fail and Delphix will raise an alert suggesting it does not have the correct sudo permissions. Restricting the execution of "sudo -l" by setting "listpw=always" in the "/etc/sudoers" file **when the Delphix operating system user is configured to use public key authentication** will cause the Delphix operating system user to be prompted for a password which will fail certain Delphix actions. Use a less restrictive setting for listpw than "always" when the Delphix operating system user is using public-key authentication.

Oracle Mount options for RAC

AIX	<code>cio,rw,bg,hard,nointr,timeo=600,proto=tcp,nosuid, noac</code>
HPUX	<code>rw,bg,hard,rsiz=1048576,wsiz=1048576,nointr,timeo=600,proto=tcp,nosuid,forcedirectio,noac</code>
Solaris	<code>rrw,bg,hard,rsiz=1048576,wsiz=1048576,nointr,pr oto=tcp,nosuid,forcedirectio,noac</code>
For the above platforms, depending on NFS version used, additional options <code>vers=3</code> or <code>vers=4.x</code> is added (x varies depending on what that platform supports. e.g. <code>vers=4</code> or <code>vers=4.1</code>)	
Linux (NFSv3)	<code>-t nfs rw,bg,hard,rsiz=1048576,wsiz=1048576,nointr,timeo=600,tcp,nosuid,sec=sys,vers=3,actimeo=0</code>
Linux (NFSv4)	<code>-t nfs4 rw,bg,hard,rsiz=1048576,wsiz=1048576,nointr,timeo=600,tcp,nosuid,sec=sys,actimeo=0</code>

- ⚠** 1. For both Single instance and RAC, "port=2049" option is added for all platforms.
2. For AIX, `rsiz=<value>`, `wsiz=<value>` options are added depending on the value returned by `"/usr/sbin/nfso -o nfs_max_read_size"` and `"/usr/sbin/nfso -o nfs_max_write_size"` commands.

Oracle unmount options

"-f" is used for all platforms. For Linux, "-lf" is used.

Oracle Mount options for Single Instance

AIX	<code>cio,rw,bg,hard,intr,timeo=600,proto=tcp,nosuid</code>
-----	---

HPUX	<code>rw,bg,hard,rsize=1048576,wsiz=1048576,nointr,timout=600,proto=tcp,nosuid,forcedirectio</code>
Solaris	<code>rw,bg,hard,rsize=1048576,wsiz=1048576,nointr,pr oto=tcp,nosuid,forcedirectio</code>
For the above platforms, depending on NFS version used, additional options <code>vers=3</code> or <code>vers=4.x</code> is added (x varies depending on what that platform supports. e.g. <code>vers=4</code> or <code>vers=4.1</code>)	
Linux (NFSv3)	<code>-t nfs rw,bg,hard,rsize=1048576,wsiz=1048576,nointr,timout=600,tcp,nosuid,sec=sys,vers=3</code>
Linux (NFSv4)	<code>-t nfs4 rw,bg,hard,rsize=1048576,wsiz=1048576,nointr,timout=600,tcp,nosuid,sec=sys</code>

Related Topics

- [Introduction to Privilege Elevation Profiles](#)
- [Sudo File Configuration Examples for Oracle Environments](#)

Sudo File Configuration Examples for Oracle Environments

This topic provides a sample `sudo` file privilege configurations for using the Delphix Engine with various operating systems and the Oracle RDBMS.

Requiretty settings

Delphix requires that the `requiretty` setting be disabled for all Delphix users with `sudo` privileges.

Configuring `sudo` Access on Solaris SPARC for Oracle Source and Target Environments

Sudo access to `pargs` on the Solaris operating system is required for the detection of listeners with non-standard configurations on both source and target environments. Super-user access level is needed to determine the `TNS_ADMIN` environment variable of the user running the listener (typically **oracle**, the installation owner). From `TNS_ADMIN`, the Delphix OS user **delphix_os** can derive connection parameters.

Example: Solaris /etc/sudoers entries for a Delphix Source

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/usr/bin/pargs
```

On a Solaris target, `sudo` access to `mount` and `umount` is also required.

Example: Solaris /etc/sudoers entries for a Delphix Target

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

User_Alias   DELPHIX_USER=delphix_os

Cmd_Alias   DELPHIX_CMDS= \
/usr/bin/mount, \
/usr/bin/umount, \
/usr/bin/mkdir, \
/usr/bin/rmdir, \
/usr/bin/pargs

DELPHIX_USER ALL=(ALL) NOPASSWD: DELPHIX_CMDS
```

Configuring `sudo` Access on Linux for Oracle Source and Target Environments

Sudo access to `ps` on the Linux operating system is required for the detection of listeners with non-standard configurations on both source and target environments. Super-user access level is needed to determine the `TNS_ADMIN` environment variable of the user running the listener (typically **oracle**, the installation owner). From `TNS_ADMIN`, the Delphix OS user **delphix_os** can derive connection parameters.

Example: Linux `/etc/sudoers` entries for a Delphix Source

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/bin/ps
```

On a Linux target, sudo access to `mount` and `umount` is also required.

Example: Linux `/etc/sudoers` file for a Delphix Target

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
/bin/mount, /bin/umount, /bin/mkdir, /bin/rmdir, /bin/ps
```

Configuring `sudo` Access on AIX for Oracle Source and Target Environments

Sudo access to `ps` on the AIX operating system is required for the detection of listeners with non-standard configurations on both source and target environments. Super-user access level is needed to determine the `TNS_ADMIN` environment variable of the user running the listener (typically **oracle**, the installation owner). From `TNS_ADMIN`, the Delphix OS user **delphix_os** can derive connection parameters.

Example: AIX `/etc/sudoers` entries for a Delphix Source

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/bin/ps
```

In addition to sudo access to the `mount`, `umount`, and `ps` commands on AIX target hosts, Delphix also requires `sudo` access to `nfso`. This is required on target hosts for the Delphix Engine to monitor the NFS read-write sizes configured on the AIX system. Super-user access level is needed to run the `nfso` command.

Example: AIX `/etc/sudoers` File for a Delphix Target

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
```

```

/usr/sbin/mount, \
/usr/sbin/umount, \
/usr/sbin/mkdir, \
/usr/sbin/rmdir, \
/usr/sbin/nfso, \
/usr/bin/ps

```

Configuring `sudo` Access on HP-UX for Oracle Source and Target Environments

No `sudo` privileges are required on source environments running HP-UX. The HP-UX OS does not allow the **delphix_os** user to determine the `TNS_ADMIN` environment variable setting for the **oracle** user. This means that the Delphix Engine cannot auto-discover non-standard listener configurations with non-default `TNS_ADMIN` settings.

On the HP-UX target, `sudo` access to `mount` and `umount` is required as with other operating systems.

Example: HP-UX `/etc/sudoers` file for a Delphix Target

Example: AIX `/etc/sudoers` File for a Delphix Target

```

Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/sbin/mount, /sbin/umount, /sbin/mkdir, /sbin/rmdir

```

Examples of Limiting `sudo` Access for the Delphix OS User

In situations where security requirements prohibit giving the Delphix user root privileges to mount, unmount, make a directory, and remove directory on the global level, it is possible to configure the `sudoers` file to provide these privileges only on specific mount points or from specific Delphix Engines, as shown in these two examples.

- i** The Delphix Engine tests its ability to run the `mount` command using `sudo` on the target environment by issuing the `sudo mount` command with no arguments. Many of the examples shown in this topic do not allow that. This causes a warning during environment discovery and monitoring but otherwise does not cause a problem. If your VDB operations succeed, it is safe to ignore this warning. Similarly, the `ps` or `pargs` command is used for target environment operations such as initial discovery and refresh. The most restrictive `sudo` setups might not allow the commands `ps` (`pargs`). Delphix can still function without these privileges, although auto-discovery may not work. However, some users configure the security on the target environments to monitor `sudo` failures and lockout the offending account after some threshold. In those situations, the failure of the `sudo` commands might cause the **delphix_os** account to become locked. One workaround for this situation is to increase the threshold for locking out the user account. Another option is to modify `/etc/sudoers` to permit the **delphix_os** user to run `ps` (`pargs`), and `mount` command without parameters.

- ⚠** Note that the following examples are for illustrative purposes and the `sudo` file configuration options are subject to change.

Example 1

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/oracle`.

Note that wildcards are allowed for the options on `mount` and `umount` because those commands expect a fixed number of arguments after the options. The option wildcard on the `mount` command also makes it possible to specify the file-system being mounted from the Delphix Engine.

Example /etc/sudoers File Configuration on the Target Environment for sudo Privileges on the VDB Mount Directory Only (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount * /oracle/*, \
/bin/umount * /oracle/*, \
/bin/umount /oracle/*, \
/bin/mkdir -p /oracle/*, \
/bin/mkdir -p -m 755 /oracle/*, \
/bin/mkdir /oracle/*, \
/bin/rmdir /oracle/*, \
/bin/ps
```

Example /etc/sudoers File Configuration on the Source Environment to grant Super-User privileges when running PS

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: /bin/ps
```

Example 2

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/oracle`, restricts the mount commands to a specific Delphix Engine hostname and IP, and does not allow user-specified options for the `umount` command.

This configuration is more secure, but there is a tradeoff with deployment simplicity. This approach would require a different sudo configuration for targets configured for different Delphix Engines.

A Second Example of Configuring the /etc/sudoers File on the Target Environment for Privileges on the VDB Mount Directory Only (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount <delphix-server-name>* /oracle/*, \
/bin/mount * <delphix-server-name>* /oracle/*, \
/bin/mount <delphix-server-ip>* /oracle/*, \
```

```
/bin/mount *          <delphix-server-ip>* /oracle/*, \  
/bin/mount "", \  
/bin/umount          /oracle/*, \  
/bin/umount *        /oracle/*, \  
/bin/mkdir, \  
/bin/rmdir, \  
/bin/ps
```

Managing Oracle Environments and Hosts

This section describes the attributes of Oracle-specific environments such as RAC Cluster Users, Listeners, and RAC configurations and covers the following topics:

- [Managing Environment Users](#)
- [Environment Attributes for Hosts with Oracle](#)
- [Adding an Oracle Single Instance or RAC Environment](#)
- [Adding Oracle Standalone and RAC Targets in Kerberos](#)
- [Adding a Database Installation Home to an Oracle Environment](#)
- [Adding a Database to an Oracle Environment](#)
- [Adding a Listener to an Oracle Environment](#)
- [Linking an Oracle RAC Source Database Using a Standalone Host Environment](#)
- [Listener and JDBC Verification](#)
- [Providing Your Own Oracle Java](#)
- [Changing the Hostname or IP Address for Oracle Source and Target Environments](#)
- [How to Change ORACLE_HOME](#)
- [CLI Support for Partial-Full Backup in Oracle](#)

Managing Environment Users

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click on the existing **environment name** you want to modify.
5. In the Details tab under **Basic Information**, click the **Plus** icon next to Environment Users to add a user.
6. Enter the **Username** and **Password** for the OS user in that environment.
7. Select **Add** to save the new user.
8. To change the primary user for this environment, click the **Pencil** icon next to **Environment Users**. Only the primary user will be used for environment discovery.
9. To delete a user, click the **Trash** icon next to their username.

Environment Attributes for Hosts with Oracle

This topic describes the attributes of Oracle-specific environments such as RAC Cluster Users, Listeners, and RAC configurations. You can manage these settings by the procedure described below.

1. Navigate to Manage > Environments.
2. In the Environments panel, click the name of an environment to view its attributes.
3. Next to Attributes, click the Pencil icon to edit an attribute.
 - For Oracle RAC environments select your node then click the Pencil icon.

Common Environment Attributes

Attribute	Description
Environment Users	The users for that environment. These are the users who have permission to ssh into an environment or access the environment through the Delphix Connector. For more information on the environment user requirements, see the Requirements topics for specific data platforms.
Host Address	The IP address of the environment host.
SSH Port	The SSH port number used for the environment.
Login Type	Select from: <ul style="list-style-type: none"> • Username and Password • Username and Public Key • Password Vault
Toolkit Path	The plugin path for the environment. Delphix uses the Toolkit path (Toolkit is a synonym for Plugin) for logs, mount points, and keeping lib files.
Notes	Any other information you want to add about the environment.

Oracle Environment Attributes

Attribute	Description
Environment Name (RAC)	The Environment Name field under Attributes is used to provide the name of the environment host in the case of cluster environments. This field defaults to the IP address of the host unless you specify another name.

Cluster Name (RAC)	The name of the RAC cluster as retrieved during auto-discovery from the Oracle Clusterware Repository (OCR).
Cluster Home (RAC)	This is the full pathname of the directory in which the cluster or grid ORACLE_HOME resides, which contains the executables for the Oracle Clusterware. Besides containing the software for the Oracle CRS (Cluster Ready Services), this directory also contains the binaries for Oracle ASM (Automatic Storage Management).
Version	The version of the RAC cluster as retrieved during auto-discovery from the Oracle Clusterware Repository (OCR).
Cluster User (RAC)	The OS user account with read-execute permission to access the ORACLE_HOME for the cluster/grid, as well as the ORACLE_HOME for the database.
Host Address (RAC)	Public IP hostname or IP address for the specific node (or host) within the RAC cluster.
NFS Address	In the event that the network link to which “public hostname” of the RAC nodes reside is to be used only for command-and-control traffic (i.e. SSH, JDBC, etc), and not for network-attached storage, then this field is where the comma-separated list of IP addresses for the alternate network dedicated to NFS for network-attached storage.
Virtual IP (RAC)	Oracle RAC requires setting up virtual IPs to manage failover. This field informs Delphix which IP addresses to use when a failure is detected. Click the + to add another virtual IP domain and IP address. For more details view Overview of Virtual IP Addresses .
Listeners	The listener is used to connect incoming client requests to the database. See Adding a Listener to an Oracle Environment for more information.
Remote Listener	An Oracle TNS name resolves to an address or address list of Oracle Net remote listeners. This is the Oracle initialization parameter REMOTE_LISTENER, which must be the same on all RAC cluster instances. Click the + to add a remote listener.
SCAN	Single Client Access Name is used to allow clients to access cluster databases.
SCAN Listener	Listener used with SCAN to establish client connections to the database instances in the RAC cluster.

Related Topics

- [Adding a Listener to an Oracle Environment](#)
- [Requirements for Oracle Hosts and Databases](#)

Adding an Oracle Single Instance or RAC Environment

This topic describes how to add a new Oracle or Oracle RAC environment.

Prerequisites

- See the topics [Requirements for Oracle Hosts and Databases](#).
- There can be one Oracle unique database name (DB_UNIQUE_NAME) per Delphix Engine. For example, if you provision a VDB with a database unique name "ABC" and later try to add an environment that has a source database that also has a database unique name of "ABC", errors will occur.

Procedure

1. Login to the **Delphix Management** application.
 2. Click **Manage**.
 3. Select **Environments**.
 4. Click the **Actions (...)** menu next to **Environments** and select **Add Environment**.
 5. In the **Host and Server** tab, select **Unix/Linux**.
 6. Select **Standalone Host** or **Oracle Cluster**, depending on the type of environment you are adding.
 7. Click **Next**.
 8. For standalone Oracle environments:
Enter the **Environment name**
Enter **Host IP Address**
 9. For cluster Oracle RAC environments:
Enter the **Environment name**
Enter **Cluster Home**
Enter **Node Address**
 10. For NFS Addresses (Optional):
Enter one or more comma-separated **IP Address/Hostname**
In the case of the Oracle RAC environment, ensure that the NFS Address list includes IP Addresses from all the cluster nodes. NFS address given at the time of environment creation is applied to all the discovered RAC nodes. If some nodes in the cluster need a different IP Address list, users can edit the specific host and update the NFS Address after the environment has been created.
-  If specified, Delphix Engine only allows NFS requests (mount, etc) originated from IP Addresses specified for the host.
11. If server authentication for remote host communication or engine to host throughput tests is desired, make sure the appropriate config is set. For more details refer to [Configuring Network Security Settings](#). You will need to create a JKS or PKCS#12 keystore on the remote host with the full CA chain of the DSP key in the keystore. By default, the key will just be signed by the Delphix CA, but you can replace the DSP key if you wish. Refer to [KeyStore Settings](#) for more details.
 12. If client authentication for remote host communication or engine to host throughput tests is also desired, make sure the appropriate config is set. For more details refer to [Configuring Network Security Settings](#). You will need to create another JKS or PKCS#12 keystore on the remote host with the desired key pair. Make sure the created keystore has permissions such that it is readable by all environment users. Then, add the full CA chain of the remote host's key pair to the TrustStore on the engine. For more details, refer to [TrustStore Settings](#).
 13. Enter the SSH port. The default value is 22.
 14. Select a Login Type:
 - Username and Password - enter the OS username and password, or
 - Username and Public Key - enter the OS username, or
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

15. For **Password Login**, click **Verify Credentials** to test the username and password.
16. Enter **Toolkit Path**.
The toolkit directory stores scripts used for Delphix Engine operations, and should have a persistent working directory rather than a temporary one. The toolkit directory will have a separate subdirectory for each database instance and must have 0770 permissions. At least 1.6 GB of storage is needed at the time of setting up the environment and at least 500MB of free space is required to allow refreshes and maintenance of the toolkit, especially during upgrades.
17. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
18. In the Java Development Kit, tab enter the absolute path to your Oracle JDK and click **Next**. For more information, see [OpenJDK on Delphix Engine](#). Adding an Oracle single instance or RAC environment Procedure
19. Click **Submit**.

Post-Requisites

After you create the environment, you can view information about it:

1. Click **Manage**.
2. Select **Environments**.
3. Select the environment name.

Related Topics

- [Configuring Network Security Settings](#)
- [KeyStore Settings](#)
- [TrustStore Settings](#)
- [Providing Your Own Oracle Java](#)

Adding Oracle Standalone and RAC Targets in Kerberos

This topic describes how to add an Oracle standalone and RAC targets in Kerberos.

Prerequisites

- Configure Kerberos credentials through the System Setup interface.
- Kerberized Oracle host operations are only supported in Delphix Engine release 5.2.3.0 and later.
- You must define only one principle in the Kerberos configuration and this principle must be used for all Kerberized host environments.

Procedure

When adding a new environment or editing the configuration of an existing environment, it is a general Kerberos authentication requirement that you configure the host address using a fully qualified domain name (FQDN). The login type of Kerberos authentication is available after you apply the Kerberos configuration in the System Setup interface.

In the example below, the Kerberos principal `krbuser` has previously been configured using System Setup, so this is automatically populated when the radio button is selected.

Perform the following steps to add an Oracle standalone and RAC targets in Kerberos.


1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Actions (...)** menu next to **Environments** and select **Add Environment**.
5. In the **Host and Server** tab, select **Unix/Linux** and click **Next**.
6. In the **Environment Settings** tab, do the following:
 - a. Enter an environment name and the host address.
 - b. Under **Login Type**, select the **Kerberos Authentication** radio button.

 The Kerberos principal `krbuser` has previously been configured using the System Setup interface, so this is automatically populated when the **Kerberos Authentication** radio button is selected.

- c. Click **Next**.
- d. The **Summary** tab enables you to review your configurations. Review the configuration and click **Submit**.

Configuring Oracle RAC using the GUI


As the discovery logic attempts to add other nodes via IP address (based on cluster interrogation results which return the additional node IP addresses), automatic discovery of additional nodes will initially fail, causing Kerberos authentication to fail with "Could not log into <IP_address> with the provided username and password". It is therefore also expected to see one or more failed actions under the Discovery job indicating "Add discovered cluster node" and "Add Oracle cluster node <node name>" were unsuccessful.



As such, all nodes beyond the first must be added manually. The cluster can be added initially using the GUI or CLI as desired to discover the first node, configured via FQDN. Each additional node is added manually via CLI or GUI.

To add all the additional nodes via GUI, perform the following steps:

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the cluster node to view the basic information of the respective node.
5. In the **Cluster Node** widget, click the plus icon to add a node. In this example, **bbrac31** was discovered normally, and **bbrac32.d.delphix.com** was manually added after the Environment Discovery process was completed.



6. Click the new node in the Cluster Nodes list.
7. Click the plus icon next to the **Virtual IPs** and add the virtual IP name and address for the node.

8. Click the plus icon next to the **Listeners** and add the name and protocol address of a listener on the node. Protocol address should be in the following format. This example uses bbrac15-vip.delphix.com as the host virtual IP address.

```
(ADDRESS=(PROTOCOL=tcp)(HOST=bbrac15-vip.delphix.com)(PORT=1521))
```

Adding an Oracle RAC Cluster using the CLI

You can use the CLI to create the cluster, specifying one of the nodes for discovery. Make sure to set the credential type to KerberosCredential. This procedure uses the following node names bbrac14, bbrac15, and bbrac16 for example purposes. "bbrac14" will be the primary node used for discovery. Make sure that you use an FQDN for the host address (it is a must Kerberos requirement for authentication):

```
delphix.engine> /environment create
delphix.engine environment create *> set type=OracleClusterCreateParameters
delphix.engine environment create *> set cluster.name=bbrac1416
delphix.engine environment create *> set cluster.crsClusterHome=/u01/app/11.2.0/grid
delphix.engine environment create *> set
primaryUser.credential.type=KerberosCredential
delphix.engine environment create *> edit nodes
delphix.engine environment create nodes *> add
delphix.engine environment create nodes 0 *> set name=bbrac14.delphix.com
delphix.engine environment create nodes 0 *> set
hostParameters.host.address=bbrac14.delphix.com
delphix.engine environment create nodes 0 *> set hostParameters.host.toolkitPath=/
work
delphix.engine environment create nodes 0 *> commit
```

Now, use the above procedure described in the **Configuring Oracle RAC using the GUI** section to manually add the node.

Adding a Database Installation Home to an Oracle Environment

When you add an environment with the Delphix Admin application, all database installation homes on it are automatically discovered. However, if a database installation home is not automatically discovered, you can add it manually to the environment.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select an Environment.
5. Click the **Databases** tab.
6. Click the **Add Dataset Home** button.
7. Enter the **Installation Home**.
8. Enter the **Version** of the Installation Home.
9. Enter the **Oracle Base** of the Installation Home.
10. Enter the **Bits** of the Oracle Home.
11. When finished, click **Add**.

Troubleshooting

If the environment user has oinstall permissions, Delphix will be able to discover the **Version**, **Oracle Base**, and **Bits**. If any of these fields are found to be different than those provided by the user, a fault will be raised on the repository.

Oracle Version

The version can be found in the comps.xml file in \$ORACLE_HOME/inventory/ContentsXML/comps.xml.

Example of Oracle Home with version 12.1.0.2:

```
<COMP NAME="oracle.server" VER="12.1.0.2.0" BUILD_NUMBER="0" REP_VER="0.0.0.0"
  RELEASE="Production" INV_LOC="Components/oracle.server/12.1.0.2.0/1/" LANGS="en"
  XML_INV_LOC="Components21/oracle.server/12.1.0.2.0/" ACT_INST_VER="12.1.0.2.0"
  DEINST_VER="11.2.0.0.0" INSTALL_TIME="2016.Apr.14 12:42:09 PDT" INST_LOC="/u01/app/
  oracle/product/12.1.0/dbhome_1/oracle.server">
```

Oracle Base

This can be found as a property in \$ORACLE_HOME/inventory/ContentsXML/oraclehomeproperties.xml.

Example of Oracle Home with Oracle Base "/u03/app/ora11202":

```
<PROPERTY NAME="ORACLE_BASE" VAL="/u03/app/ora11202"/>
```

Bits

This can be found by running

```
file $ORACLE_HOME/bin/rman
```

The output will indicate if the Oracle Home is 32 bit or 64 bit.

Related Topics

- [Environment Attributes for Hosts with Oracle](#)
- [Requirements for Oracle Hosts and Databases](#)

Adding a Database to an Oracle Environment

Prerequisites

- Make sure your source database meets the requirements described in [Requirements for Oracle Hosts and Databases](#).
- Before adding a database, the installation home of the database must exist in the environment. If the installation home does not exist in the environment, follow the steps in [Adding a Database Installation Home to an Oracle Environment](#).

Procedure

1. Log in to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select an Environment.
5. Click the **Databases** tab.
6. Choose the installation home where the database is installed.
7. Click the **Plus** icon.



8. In the Add Database dialog enter the **Database Path** and **Port**.

9. When finished, click **Add**.

Related Topics

- [Requirements for Oracle Hosts and Databases](#)
- [Adding a Database Installation Home to an Oracle Environment](#)

Adding a Listener to an Oracle Environment

This topic describes how to add listeners for an Oracle environment.

When you add an environment with the Delphix Management application, all listeners that are running on it are automatically discovered. However, if a listener is not automatically discovered, you can add it manually to the environment.

Procedure

1. Log into the **Delphix Management** application.
2. Select **Manage > Environments**.
3. Click on the name of an environment to view its basic information.
4. In the Details tab next to Listeners, click the **Pencil** icon to edit the list of listeners.

Listeners 

Name
LISTENER

Protocol addresses
(ADDRESS=(PROTOCOL=tcp)(HOST=10.43.8.93)(PORT=1521))
(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521))

- In the Listeners panel click the **Plus** icon.

The screenshot shows the 'Listeners' panel in the Delphix Engine interface. It contains three listener entries. The first two entries are in an edit state, each with a 'Name' input field and a 'Protocol addresses' input field containing the text '(ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=1521))'. A plus icon is visible next to the 'Protocol addresses' label for each. The third entry is in a view state, showing the name 'LISTENER' and two protocol addresses: '(ADDRESS=(PROTOCOL=tcp)(HOST=10.43.8.93)(PORT=1521))' and '(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC1521))'. A plus icon is also visible in the top right corner of the panel.

- Enter a **Name** for the new listener, and a **Protocol Address** of the listener. The Delphix Engine currently supports TCP and IPC protocol addresses. An example TCP protocol address is (ADDRESS=(PROTOCOL=TCP)(HOST=10.43.17.92)(PORT=1521)) and an example IPC protocol address is (ADDRESS=(PROTOCOL=IPC)(KEY=DELPHIX))
- Click the **Plus** icon next to **Protocol Addresses** to enter additional protocol addresses.
- Click the **Check** icon to save your changes.

Related Topics

- [Requirements for Oracle Hosts and Databases](#)
- [Listener and JDBC Verification](#)

Linking an Oracle RAC Source Database Using a Standalone Host Environment

The standard way to link an Oracle RAC database as a dSource in the Delphix workflow is to provide the name or IP of one of the cluster nodes and then discovery will find the other cluster members, listeners, and databases. This is done by adding an 'Oracle Cluster' type Environment in the Add Environment wizard. The cluster nodes will be added to the Environment using the hostname/IP returned by the `'olsnodes'` command and in most cases, this will be sufficient. There may, however, be cases when this is undesirable, such as wanting to run snapsync/logsync traffic over a second network. Using an 'Oracle Cluster' type Environment currently stores a single IP for each cluster node and this IP will be the one associated with the node name returned by the `'olsnodes'` command. Using a Standalone Environment is one way to work around this situation by defining the Environment based on the IP on the network snapsync/logsync should run over. The Delphix Engine should have a NIC connected to this network also. When both Delphix Engine and source servers have NICs on this second network there is no need for static routes to be configured on the Delphix Engine or the source server.

1. Ensure that OS and DB users for Delphix have been configured and [meet requirements](#). Create the OS user on each node using the exact same settings. Always use the hostchecker to verify the source and target environments are ready for use with Delphix.
Even though a single node will be used for linking, it may be desirable to unlink and relink the source to a different node in the future, so make sure the Delphix OS account is configured identically on all nodes.
2. Create a toolkit directory in the same location on each node and ensure permissions are set to 0770.
3. Identify the local IP address on the second network of the node which will be used for linking. Use the `'oifcfg'` and `'ifconfig'` commands to find this information. Then verify by pinging the Delphix Engine using its IP on the second network. Don't worry about the output from `'oifcfg'` with respect to `PUBLIC, PRIVATE` or `UNKNOWN` values for the interfaces. These are not actual CRS usage, but rather based on the IP address range (RFC1918) - in the image they all show as `PRIVATE`.

```
[alex:syrinx_3]/home/oracle => $CRS_HOME/bin/oifcfg iflist -p -n
eth0 10.0.0.0 PRIVATE 255.255.0.0
eth1 192.168.1.0 PRIVATE 255.255.255.0
eth1 169.254.0.0 UNKNOWN 255.255.0.0
eth2 172.16.0.0 PRIVATE 255.240.0.0
[alex:syrinx_3]/home/oracle => ifconfig eth2
eth2      Link encap:Ethernet  HWaddr 08:00:27:6A:A9:2C
          inet addr:172.16.16.100  Bcast:172.31.255.255  Mask:255.240.0.0
          inet6 addr: fe80::a00:27ff:feba:a92c:b4 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:47584 errors:0 dropped:0 overruns:0 frame:0
          TX packets:132277 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:5026787 (4.7 MiB)  TX bytes:168273350 (160.4 MiB)

[alex:syrinx_3]/home/oracle => ping 172.16.16.85
PING 172.16.16.85 (172.16.16.85) 56(84) bytes of data.
64 bytes from 172.16.16.85: icmp_seq=1 ttl=255 time=0.301 ms
64 bytes from 172.16.16.85: icmp_seq=2 ttl=255 time=0.534 ms
64 bytes from 172.16.16.85: icmp_seq=3 ttl=255 time=0.251 ms
64 bytes from 172.16.16.85: icmp_seq=4 ttl=255 time=0.483 ms
64 bytes from 172.16.16.85: icmp_seq=5 ttl=255 time=0.358 ms
^C
--- 172.16.16.85 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4372ms
rtt min/avg/max/mdev = 0.251/0.385/0.534/0.108 ms
[alex:syrinx_3]/home/oracle => _
```

4. Add a 'Standalone Host' type environment using the chosen cluster node's second network IP.
5. When discovery is finished, the Oracle homes on the cluster node should be discovered along with the listeners that run on that node. Once that is done add the database that will eventually be a dSource to by providing the db_unique_name, db_name, and the instance name that runs on the node that was added.
6. Then add a jdbc connect string using the node VIP and a service name.
7. Now click on the add dSource link. Follow the steps in [Adding a dSource](#) making sure to set the Snapshot controlfile location in RMAN to a location that's visible to all cluster nodes, or snapsync will fail. An ASM diskgroup is a suitable location.

Even if the Snapshot controlfile is not located in a shared location, it is still possible to link the source - the first snapsync will fail with a fault that relates to the Snapshot controlfile location, but subsequent Snapsync operations should succeed.

JDBC connections will go over the public network and snapsync / logsync / environment monitor traffic will go over the second network.

Related Topics

- [Linking Data Sources with Oracle](#)
- [Oracle Support and Requirements](#)

Listener and JDBC Verification

As with any database, a VDB needs to have a listener in order for external connections to be made. The Delphix Engine also uses a JDBC connection string in order to connect to a source or target database. If the engine does not have the proper listener or JDBC connect string defined, then connection errors can result.


Verifying the Listener Configuration

1. Verify that a listener is running on the source or target environment that you are investigating.
2. Login to the Delphix Management application.
3. Click **Manage**.
4. Select **Environments**.
5. Click on the **environment in** which you are troubleshooting.
6. On the **Configuration** tab, locate the **Listeners** section. Verify that the listener for the source or target system is listed there.
7. If a listener is not listed, it may be due to insufficient privileges on the part of the environment user that the engine is using. Verify that the proper sudo permissions have been granted for the user, or adjust them as necessary.
8. If you need to add another listener, you can do so manually.
 - a. Click the plus (+) sign.
 - b. Enter the appropriate values using the above image as a reference.

Verifying the JDBC Connection String

Each source database and VDB has a connection string defined. If any parameters have changed, you may need to adjust these connection strings.

1. Verify that a listener is running on the source or target environment you are investigating.
2. Login to the Delphix Management application.
3. Click **Manage**.
4. Select **Environments**.
5. Click the environment in which you are troubleshooting.
6. Select the **Databases** tab.
7. The source database or VDB which you are investigating will display the JDBC connection string being used for the given database.



8. To verify that the connection string works, click the **checkmark** to the right of the connection string. You will then see the username and password text boxes.

9. Enter the Oracle **username** and **password** used by the Delphix Engine.

Verify Connection String ✕

JDBC Connection String
jdbc:oracle:thin:@(DESCRIPTION=(ENABLE=broken)(ADDRESS=(PROTOCOL=tcp)
(HOST=10.110.222.165)(PORT=1521))(CONNECT_DATA=(UR=A)
(SID=CDOMLOSRB5A3)))

Database Username

Database Password

If no errors are encountered and the username/password dialogs simply go away, then a successful connection was made to the database.

If errors are encountered, you must investigate them on a case-by-case basis, just as you would do with any connection errors to an Oracle database.


Related Topics

- [Listener and JDBC Verification](#)

Adding JDBC Connection String


If you need to define an additional JDBC connection string for any reason, take the following steps:

1. Follow the same steps as listed in [Verifying the JDBC Connection String](#) to reach the JDBC connection string definition(s).
2. Click the plus (+) sign to the right of **Connection Strings** to define an additional connection string.



Add Connection String Icon

In the screenshot below, the user is adding a connection string for the “test2” service on port 1530 instead of the default 1521.



3. Click the **Add** to save the changes.
4. Follow the remaining steps in [Verifying the JDBC Connection String](#) to validate your newly added connection string.

Related Topics

- [Verifying the JDBC Connection String](#)

Verifying the JDBC Connection String

Each source database and VDB has a connection string defined. If any parameters have changed, you may need to adjust these connection strings.

1. Verify that a listener is running on the source or target environment you are investigating.
2. Login to the Delphix Management application.
3. Click **Manage**.
4. Select **Environments**.
5. Click the environment in which you are troubleshooting.
6. Select the **Databases** tab.
7. The source database or VDB which you are investigating will display the JDBC connection string being used for the given database.
8. To verify that the connection string works, click the **checkmark** to the right of the connection string. You will then see the username and password text boxes.
9. Enter the Oracle **username** and **password** used by the Delphix Engine.

Verify Connection String ✕

JDBC Connection String
 jdbc:oracle:thin:@(DESCRIPTION=(ENABLE=broken)(ADDRESS=(PROTOCOL=tcp)(HOST=10.110.222.165)(PORT=1521))(CONNECT_DATA=(UR=A)(SID=CDOMLOSRB5A3)))

Database Username

Database Password

Database Username and Password

If no errors are encountered and the username/password dialogs simply go away, then a successful connection was made to the database.

If errors are encountered, you must investigate them on a case-by-case basis, just as you would do with any connection errors to an Oracle database.

Related Topics

- [Adding JDBC Connection String](#)

Verifying the Listener Configuration

1. Verify that a listener is running on the source or target environment that you are investigating.
2. Login to the Delphix Management application.
3. Click **Manage**.
4. Select **Environments**.
5. Click on the **environment in** which you are troubleshooting.
6. On the **Configuration** tab, locate the **Listeners** section. Verify that the listener for the source or target system is listed there.
7. If a listener is not listed, it may be due to insufficient privileges on the part of the environment user that the engine is using. Verify that the proper sudo permissions have been granted for the user, or adjust them as necessary.
8. If you need to add another listener, you can do so manually.
 - a. Click the plus (+) sign.
 - b. Enter the appropriate values using the above image as a reference.

Related Topics

- [Adding a Listener to an Oracle Environment](#)

Providing Your Own Oracle Java

Delphix supports the ability for users to provide their own Oracle JDK for each host. This feature is intended for customers who have a Java license with Oracle and want to use a more recent version than the one that comes with Delphix. Note that the JDK is provided per node or host, not per cluster or environment. Initially, when creating a cluster all hosts must have Java at the same absolute path. However, once the cluster has been discovered in Delphix each host's path can be adjusted.

For details about which versions of Oracle Java are supported see the Java Support Policy section below.

⚠ All Delphix environment users on the host require read and execute permissions on the provided JDK, its subfolders, and files.

Adding an Oracle JDK

By default, Delphix comes with the OpenJDK. To modify the JDK, follow the below steps:

1. log in to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Actions (...)** menu next to **Environments** and select **Add Environment**.
5. In the **Environment Settings** tab, select the **Provide my own JDK** checkbox, and click **Next**.

This action will remove the previous built-in JDK and will initiate an Environment refresh operation after the path is changed.

6. In the **Java Development Kit** tab, the currently selected JDK kit will be shown (default is OpenJDK). Provide an absolute path to the JDK root and click **Next**.

Add Environment

- Host and Server
- Environment Settings
- Java Development Kit**
- Summary

Java Development Kit

The current Java Development Kit that is installed:

Default JDK

Java Development Kit (JDK) Path ⓘ

 Select the Reset to Default to reset your JDK to the OpenSource JDK (default).

7. Verify the configured JDK and environment settings and click **Submit**.

Updating a JDK Path

For existing environments users can add or update JDK paths by following the steps below:

1. Log in to the **Delphix Management** application.
2. Select **Manage > Environments**.
3. Click on the name of an environment to view its basic information.
4. In the Details tab next to the **Java Development Kit**, click the **Pencil** icon.

Java Development Kit

The current Java Development Kit that is installed:

Default JDK

Java Development Kit (JDK) Path ⓘ

✕
✓

5. In the **Java Development Kit dialog**, provide an absolute path to the JDK root and click **Next**.

 Do not place the JDK inside the Delphix Toolkit.

6. Click the **Check** icon to save your changes.

Java Support Policy

For each release of the Delphix engine, we will support Java version(s) depending on the available Java builds provided by Oracle in that time period. The following statements apply for Java support for specific Delphix versions:

1. In accordance with our security policy, every six months, we will certify the latest update release of the currently supported LTS Oracle Java.
2. We will support at least one commercially supported LTS release per Delphix version for 1 year, at which point we will update the minimum supported version.
3. Introducing support for new LTS versions of Java will be considered and evaluated with each major version of Delphix.
4. At least 1 year before Oracle ends support for an LTS version of Java, Delphix will also end support for that version and update it to a newer LTS version.

Supported Operating Systems

This feature is supported with use for Linux, Solaris, and Windows environments, as described in the support matrix. We do not support Bring Your Own Oracle Java with AIX and HP-UX because those vendors provide us with a specific Java version to run on those hosts.

Only Java versions equal to or greater than U242 and equal to or lesser than U281 are tested and certified by Delphix.

	RHEL	SLES	Solaris	AIX	HP-UX	Windows
Oracle Java 8 u242	Supported	Supported	Supported	N/A	N/A	Supported
Oracle Java 8 u281	Supported	Supported	Supported	N/A	N/A	Supported

Delphix Toolkit Native Java Support Matrix

This matrix describes the supported versions of Java we package with the Delphix toolkit for each operating system. This is the default option if you do not use the feature to provide your own Java for each host.

	RHEL	SLES	Solaris	AIX	HP-UX	Windows
AdoptOpenJDK 8u302	Supported	Supported	Supported	N/A	N/A	Supported
IBM Java 8.0.0.620	N/A	N/A	N/A	Supported	N/A	N/A
HP Java 8.0.21	N/A	N/A	N/A	N/A	Supported	N/A

Related Topics

- [Oracle Support and Requirements](#)
- [Managing Oracle Environments and Hosts](#)

Changing the Hostname or IP Address for Oracle Source and Target Environments

Procedure

For Source Environments

1. Disable the dSource as described in [Enabling and Disabling Oracle dSources](#).
2. If the **Host Address** field contains an IP address, edit the IP address.
3. If the **Host Address** field contains a hostname, update your Domain Name Server to associate the new IP address to the hostname. The Delphix Engine will automatically detect the change within a few minutes.
4. In the **Environments** screen of the Delphix Engine, refresh the host.
5. Enable the dSource.

For VDB Target Environments

1. Disable the VDB as described in [Enabling and Disabling Oracle VDBs](#).
2. If the **Host Address** field contains an IP address, edit the IP address.
3. If the **Host Address** field contains a hostname, update your Domain Name Server to associate the new IP address to the hostname. The Delphix Engine will automatically detect the change within a few minutes.
4. In the **Environments** screen of the Delphix Engine, refresh the host.
5. Enable the VDB.

For the Delphix Engine

1. To stop running your VDB select the red **Stop** button located on the VDB **Configuration** tab.
2. Disable all dSources as described in [Enabling and Disabling Oracle dSources](#).
3. You can use either the command-line interface or the Delphix Setup application to change the IP address of the Delphix Engine.
 - a. To use the command-line interface, press **F2** and follow the instructions described in [Setting Up Network Access to the Delphix Engine](#).
 - b. To use the Delphix Setup application, go to **Delphix Management > Engine Setup** in the Delphix Management application, or click **Server Setup** in the Delphix Engine login screen.
 - i. In the **Network** panel, click **Modify**.
 - ii. Under **DNS Services**, enter the new IP address.
 - iii. Click **OK**.
4. Refresh all Environments by clicking the **Refresh** symbol on the **Environments** screen.
5. Enable all dSources as described in [Enabling and Disabling Oracle dSources](#).
6. To start all VDBs, click the **Start** button located on the VDB **Configuration** tab.

Using Custom init.ora or spfile.ora Files

If you are using custom `init.ora` or `spfile.ora` files with your Oracle VDBs, you should use the Oracle command-line interface (`sqlplus/srvctl`) to shut down any active VDBs and copy the parameter files to a backup location. Complete the steps above, then replace the files and re-start the VDB from the Oracle command line to restore your custom settings. See [Customizing VDB File Mappings](#) for more information about customizing `init.ora` and other configuration files.

Related Topics

- [Setting Up Network Access to the Delphix Engine](#)
- [Managing Data Sources and Syncing Data](#)

- [Provisioning and Managing Virtual Databases](#)
- [Customizing VDB File Mappings](#)

How to Change ORACLE_HOME

Prerequisites

- know the new ORACLE_HOME user and password
- know the VDB you want to change
- have delphix_admin privileges to the VDB

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select your VDB.
5. Select the **Configuration tab**.
6. In the upper right-hand corner of the **Source** sub-tab, click the **Actions menu (...)** and select **Upgrade**.
7. Select the new **Installation** from the drop-down menu.

The screenshot shows a dialog box titled "Upgrade Database" with a close button (X) in the top right corner. The dialog is divided into three main sections:

- Current Installation:** This section contains four fields: "Database", "Installation", "Environment", and "User". Each field has a greyed-out placeholder value.
- New Installation:** This section contains a single dropdown menu labeled "Installation" with the value "/u05/app/ora12201/product/12.2.0/dbhome_1" selected.
- Buttons:** At the bottom right of the dialog, there are two buttons: a grey "Cancel" button and a blue "Upgrade" button.

8. Click **Upgrade**.

Related Topics

- [Managing Oracle Environments](#)

CLI Support for Partial-Full Backup in Oracle

Overview

Delphix Engine version 6.0.8.0 introduces an option to perform partial full snapsync via the CLI. A full backup of specified datafiles will be performed, along with an incremental backup of the remaining files.

Prerequisites

1. Requested datafiles must exist in the database for which SnapSync is being run.
2. In case of an MT, Snapsync must be performance for the PDB, the datafiles can belong to either the PDB or its CDB.
3. Requested datafiles must be online.
4. Partial full SnapSync can not be requested for an initial SnapSync.
5. Partial full SnapSync can not be run if there is a pending SnapSync.

Command

```
ip-00-000-000-000 database> select CDOMLOSR1CEBPDB3
ip-00-000-000-000 database 'CDOMLOSR1CEBPDB3'> sync
ip-00-000-000-000 database 'CDOMLOSR1CEBPDB3' sync *> set filesForFullBackup=1,2
ip-00-000-000-000 database 'CDOMLOSR1CEBPDB3' sync *> commit
```

Linking Data Sources and Syncing Data with Oracle

Linking a data source will create a dSource object on the engine and allow Delphix to ingest data from this source. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

Once a dSource has been linked, you can begin to perform dSource operations on it such as enable, detach, delete, and more.

This section covers the following topics:

- [Linking Data Sources with Oracle](#)
- [Linking an Oracle Pluggable Database](#)
- [Data Management Settings for Oracle Data Sources](#)
- [Prepare and Upgrade a Non-MT Oracle dSource to MT](#)
- [Repository Templates for Oracle Databases](#)
- [Using the PurgeLogs Operation](#)
- [Advanced Linking](#)

Linking Data Sources with Oracle

Linking a dSource will ingest data from the source and create a dSource object on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

Once a dSource has been linked, you can begin to perform dSource operations on it such as enable, detach, delete, and more.

For an overview of all dSource related actions, please visit [Provisioning and Managing Virtual Databases](#).

1. Log in to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.
4. In the **Add dSource** wizard, select the source database with the correct environment user-specified.
5. Enter your login credentials for the source database and click **Verify Credentials**.
If you are linking a mounted standby, Click **Next**. See the topics under [Linking Oracle Physical Standby Databases](#) for more information about how the Delphix Engine uses non-SYS login credentials.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed. For more information, visit [Data Management Settings for Oracle Data Sources](#).
8. Assign existing policies to the new dSource. New policies can be created and associated later.
9. Enter any scripts that should be run using the **Hooks** page.
10. Review the dSource Configuration and Data Management information, and then click **Submit**.

Once the action to add a dSource has been submitted, the Delphix Engine will initiate a DB_Link job to create the dSource. If the *Load Immediately* option was selected in the data management page a DB_Sync job will also be executed to ingest data from the source, otherwise, this first DB_Sync job will run as per the associated SnapSync policy.

When the jobs have successfully completed, the database icon will change to a dSource icon on the Environments > Databases screen, and the dSource will be added to the list of Datasets under its assigned group.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)
- [Provisioning and Managing Virtual Databases](#)
- [Linking Oracle Physical Standby Databases](#)
- [Data Management Settings for Oracle Data Sources](#)

Linking an Oracle Pluggable Database

This topic describes how to link an Oracle 12c pluggable database to the Delphix Engine to create a dSource.

Prerequisites

- Make sure the Delphix Engine has already discovered the multitenant container database and its pluggable databases. If the container database does not exist in the environment, follow the steps in [Adding a Database to an Oracle Environment](#). If the pluggable database you want to link does not exist in the environment.
- You should have Block Change Tracking (BCT) enabled for the container database, as described in [Requirements for Oracle Hosts and Databases](#).
- The container database should be in ARCHIVELOG mode and the NOLOGGING option should be disabled, as described in [Requirements for Oracle Hosts and Databases](#).

Procedure

1. Log into the **Delphix Management** application.
2. Select **Manage > Datasets**.
3. Click **the Plus** icon and select **Add dSource**.
4. In the **Add dSource** wizard, select the source pluggable database.
If the container database is shown but the pluggable database is not, select the **container database**, enter its **database credentials**, and click **Verify Credentials**. The Delphix Engine will discover and list all pluggable databases in the container database. Select the **pluggable database** from the list. Alternatively, you can use the **Environment Management** screen to discover pluggable databases in any of the discovered container databases.
5. Enter your **login credentials** for the source database and click **Verify Credentials**.
6. Click **Next**.
7. Select a **Database Group** for the dSource.
8. Click **Next**.
9. Select an **Initial Load** option.
By default, the initial load takes place upon completion of the linking process. Alternatively, you can set the initial load to take place according to the SnapSync policy. For example, you can set the initial load to take place when the source database is not in use, or after a set of operations have taken place.
10. Select a **SnapSync** policy.
See [Advanced Data Management Settings for Oracle dSources](#) for more information.
11. Click **Next**.
12. Review the **dSource Configuration** and **Data Management** information.
13. Click **Submit**.
The Delphix Engine will initiate two jobs, **DB_Link** and **DB_Sync**, to create the dSource. You can monitor these jobs by clicking **Active Jobs** in the top menu bar, or by selecting **System > Event Viewer**. When the jobs have completed successfully, the database icon will change to a dSource icon on the **Environments > Datasets** screen, and the dSource will be added to the list of **Datasets** under its assigned group.
Link/Sync of the Multitenant Container Database
The **DB_Link** job will also link the pluggable database's multitenant container database if it has not been linked yet.

Related Topics

- [Requirements for Oracle Hosts and Databases](#)
- [Adding a Database to an Oracle Environment](#)
- [Data Management Settings for Oracle Data Sources](#)

Data Management Settings for Oracle Data Sources

Each dSource has its own data management settings, which can be configured during the linking workflow as well as in the configuration page for that dSource.

You can configure data management settings to improve overall performance and match the needs of your specific server and data environment.

The following settings are available for Oracle data sources:

Setting	Explanation
Initial Load	By default, the initial load takes place upon completion of the dSource linking wizard. Alternatively, you can set the initial load to take place according to a SnapSync policy. This might be desirable if you want the initial load to take place when the source database is under a lower load.
External Data Directory	An optional directory on the database server containing extra files associated with the database that are not included in normal SnapSync snapshots, for example, external table data files. Nothing in this directory will be modified, but the files will be loaded and provisioned along with the database.
Compression	Enable compression of backup data sent over the network. Compression is enabled by default.
Bandwidth Limit	Select the network bandwidth limit in units of MB/s between the source and the Delphix Engine. The default is 0, indicating that no bandwidth limit is enforced.
Number of Connections	Select the number of TCP connections to use between the Source and the Delphix engine. Multiple connections may improve network throughput especially over long distances and highly congested networks. The default is 1.
Encrypted Linking	Turn on encryption between the source and the Delphix Engine. Enabling encrypted linking leads to higher CPU utilization and higher throughput. Encrypted Linking is disabled by default.

Data Load Channels	<p>The channel settings determine the number of channels and data files per backup set. While these settings can be increased, you should consider potential adverse effects on the performance of database operations on the Source server</p> <ul style="list-style-type: none"> Number of Channels - set the number of RMAN channels used during SnapSync. The default is 2. Files per Channel - maximum number of data files in a backup set. The default is 5. <p>The product of files-per-channel and number of channels determines the maximum number of data files concurrently backed up by RMAN.</p>
Block Checking	<p>Enable logical block validation by RMAN. Checking is CPU intensive and will slow down SnapSync. Enabling block checking will detect logical inconsistencies while syncing with the source.</p> <p>Block checking is disabled by default.</p>
Level Backup	<p>Level backups should only be used to workaround Oracle bug 10146187 on physical standby sources. Switching from SCN to LEVEL mode will force a Level 0 backup.</p> <p>See Linking Oracle Physical Standby Databases for more information. Level Backups are disabled by default.</p>
Skip Available Space Check	<p>Skips checks for available space during the initial load in the Delphix Engine. The initial load will fail if the database needs more space than what is available in the Delphix Engine.</p> <p>This setting is only recommended after consulting with Delphix Support</p>
Double SnapSync	<p>Double SnapSync, or DoubleSync, takes two snapshots sequentially in order to improve the performance of subsequent operations, such as VDB provisioning and refresh.</p> <p>Performs the initial load without retrieving any archived logs, thereby creating a non-provisionable snapshot. At the completion of the initial load a second SnapSync that will retrieve logs will start. For more information, view the KB article.</p>
LogSync Policy Settings	<p>LogSync Policy Settings are available after the dSource has been added. On the Configuration page, you can select the following options:</p> <ul style="list-style-type: none"> Enabled - LogSync fetches log files from the source database, enabling the ability to provision a VDB from a specific point in time or, a database change number (SCN in the case of Oracle databases). LogSync must be enabled for this provisioning functionality to work. Archive Only, Archive and Online Redo - these settings determine whether LogSync fetches logs from archive storage in the source database file system, or both the file system and online redo logs. Online Redo mode is not supported for physical standby or RAC databases. It is also not supported if the online redo logs are raw devices or Oracle Automatic Storage Management devices. If LogSync detects any of the restricted cases it will automatically enter into Archive Only mode, regardless of the mode that was chosen. Additionally, Online Redo mode is not supported for mounted but not open primary databases. <p>LogSync policy settings for Oracle pluggable databases must be set at their corresponding container databases.</p>

Validated Sync	<p>Oracle validated sync is disabled by default. When enabled, validated sync is performed immediately after every subsequent SnapSync. See Enabling Validated Sync for Oracle for more information.</p> <p>Note: Validated sync for Oracle pluggable databases is not supported in this release</p>
----------------	---

Snapshot Parameters

There are also parameters you can specify when creating a snapshot of an Oracle data source. These parameters are explained in the table below.

Parameter	Explanation
Force Full Backup	The Delphix Engine will by default perform an incremental backup, if for some reason a full backup is required (most often to resolve issues with datafile corruption), this option should be used.
Double Sync	<p>Double SnapSync, or DoubleSync, takes two snapshots sequentially in order to improve the performance of subsequent operations, such as VDB provisioning and refresh.</p> <p>Performs the initial load without retrieving any archived logs, thereby creating a non-provisionable snapshot. At the completion of the initial load a second SnapSync that will retrieve logs will start. For more information, view the Using the Double Sync option for Oracle SnapSync.</p>
Do Not Resume	During the initial SnapSync, if a failure is encountered, the Delphix Engine can resume the SnapSync at a later date, this option will cause the engine to not resume, but rather to start the initial SnapSync over again.

Related Topics

- [Linking Oracle Physical Standby Databases](#)
- [Enabling Validated Sync for Oracle](#)

Prepare and Upgrade a Non-MT Oracle dSource to MT

Overview

This article describes how Delphix works with production databases that are converted from non-multitenant to multitenant, and the workflow for this operation. Once a dSource is converted to a multitenant PDB, it will be able to share storage blocks with its non-multitenant predecessor, and Delphix will only store the incremental changes to the database. This operation is known as STConvertedToPDBAttach and is only available via CLI.

Environment Requirements

1. Delphix Engine 6.0.10.0 or later.
2. Source host with a 11g or 12c non-multitenant source database.
3. PDB that is created as a result of converting the above non-multitenant database to multitenant. This PDB can reside on a different host.
4. Source host with a CDB that hosts the above PDB.

Workflow Steps

At the end of the steps below, there should be 2 containers:

- The detached non-MT container which contains all the non-MT snapshots.
- A new container representing the converted PDB and all its subsequent snapshots with the container name as the PDB name. By default the new MT container will have the name of the PDB, or a parameter can be set by the customer to specify its name during attach.

Use these steps to perform this functionality.

1. Link the non-multitenant Oracle 11g or 12c source database as a dSource within Delphix. Please note, 11g sources Patch 31220011 included in the July 2020 bundle is needed and XML must be installed to the DB.
2. Before performing the conversion to multitenant, detach the dSource and shut it down. Visit the [Detach/Re-Attach Oracle dSource](#) article for this step and step 7.
3. Convert the non-multitenant database to a PDB following procedures provided in the Oracle documentation. Delphix recommends that the new PDB have the same name as the original non-multitenant database. The new CDB name should be different from the original non-multitenant database.
4. Discover the CDB in Delphix Engine if not already present.
5. Refresh the environment to make sure the new PDB is discovered.
6. If the PDB name has the same name as the non-multitenant then the old non-MT container can be renamed or a parameter to give a different name to the PDB database should be used.
7. Attach converted PDB to non-multitenant dSource using CLI (attach-converted-pdb sample script).

Converted PDB Rules BEFORE SnapSync

- Attachment of a converted PDB does not perform automatic SnapSync, even if linkNow=true is used. This is to allow detachment of a converted PDB without negative effects on a non-MT container.
- If for any reason it is necessary to re-convert the PDB (after re-opening to the non-MT database to perform changes), then the converted PDB should be removed instead of detached so that the STConvertedToPDBAttach can be done again to save space.
- Once the converted PDB is removed without performing SnapSync on it, it is possible to attach the non-MT container to its non-MT database.

i **Converted PDB Detach vs. Delete BEFORE SnapSync**

- Detaching a converted PDB will make it possible to attach the non-MT container to its non-MT database.
- However, if the non-MT database is reconverted and regular attach is used to attach the PDB, there will be no space de-duplication since the detach operation reverted the changes as it was before the STConvertedToPDBAttach operation.
- Detach of a converted PDB and subsequent attach should only be performed as long as the PDB has not been re-created/re-plugged to obtain space de-duplication.
- In general a converted PDB should be removed if no SnapSync is taken.

i **Converted PDB Rules AFTER SnapSync**

- Once a SnapSync is taken of a converted PDB, the non-MT container cannot be re-attached to any other container (neither to a non-MT database or to a converted pdb).
- This means that once SnapSync is taken, the non-MT container cannot be used to attach to a converted PDB anymore.
- Even if the converted PDB is removed after taking a SnapSync, attaching to the non-MT container will not be possible as the non-MT container is marked as already converted.

AttachSource/Data Properties

These are all the possible properties that can be configured during the attach of a converted PDB. Of these, only `newPdbContainer` and `force` have special meaning for a converted PDB. The rest are 'regular' attachment options.

```
machine.name database 'DBOMSR' attachSource attachData *> ls
Properties
  type: OracleSTConvertedToPDBAttachData (*)
  backupLevelEnabled: (unset)
  bandwidthLimit: (unset)
  checkLogical: (unset)
  compressedLinkingEnabled: (unset)
  config: (required)
  encryptedLinkingEnabled: (unset)
  environmentUser: (required)
  externalFilePath: (unset)
  filesPerSet: (unset)
  force: (unset)
  linkNow: (unset)
  newPdbContainerName: (unset)
  numberOfConnections: (unset)
  operations: (unset)
  oracleFallbackCredentials: (unset)
  oracleFallbackUser: (unset)
  rmanChannels: (unset)
```

attach-converted-pdb sample script

```
cd /database
select mydatabase
attachSource
edit attachData
set type=OracleSTConvertedToPDBAttachData
set newPdbContainerName=<name for the container instead of PDB name>
set config=<converted PDB name>
edit oracleFallbackCredentials
set type>PasswordCredential
set password=<database login password>
back
set oracleFallbackUser=<database login username>
set environmentUser=<environment user>
commit
```

Repository Templates for Oracle Databases

The primary use case and motivation for repository templates are to provide the Delphix administrator with control over the Oracle database parameters used during the staging phase of the VDB provisioning process. It is useful to be able to control these configuration parameters when the physical capabilities of the staging machine, such as CPU count and memory, are smaller than the physical capabilities of the machines hosting the source database repository.

The repository template is a relationship between three entities:

- A database repository – It is the Oracle home to which the vPDB will be provisioned on the target host.
- A database container – It is the source PDB.
- A VDB configuration template – A list of database configuration parameter names and values that you can save on the Delphix Engine to use at a later time.

The VDB configuration template for the temporary CDB must contain at least the following database parameters:

- **compatible**: Required for both Non-Multitenant and Multitenant sources. Must be set with the same value as in the dSource 'compatible' parameter.
- **enable_pluggable_database**: Required for Multitenant sources only. Must be set to TRUE.

During the staging process, if you do not specify a repository template, then by default, the Delphix Engine will use the configuration parameters taken from the source database to configure the staged database. These parameters may not be appropriate, because the machine used for staging may be physically inferior to the machine hosting the source database.

Instead, the Delphix administrator can create a VDB configuration template, which would be appropriate for the physical machine hosting staging repository. (See [VDB Config Templates](#).) Then the admin can create a repository template entry that will bind together the VDB configuration template, database repository, and database container. This instructs the Delphix Engine to use configuration parameters from the VDB configuration template whenever the database container is staged on the database repository specified, instead of the parameters on the source database.

Currently, repository template relations can only be created via the command-line interface (CLI) in the repository > template.

Provisioning a VPDB using a Repository Template

Perform the following procedure to provision a VPDB using a repository template.

1. Create a VDB configuration template using the Delphix GUI called `aux_cdb_params` with the following parameters set as a bare minimum.

```
compatible = 12.1.0.2
enable_pluggable_database = true
sga_target = 1551892480
```

The parameters required may differ given a specific environment's requirements from this but this should allow the startup of the auxiliary database.

2. Create a repository template called `aux_cdb_tmpl` using the CLI.

```
delphix repository template> create
delphix repository template create *> set name=aux_cdb_tmpl
delphix repository template create *> set container=PDB12C1
```

```
delphix repository template create *> set repository=OELC9/'/u01/app/oracle/
12.1'
delphix repository template create *> set template=aux_cdb_params
delphix repository template create *> commit
```

```
delphix repository template 'pb12c1rt'> ls
Properties
  type: SourceRepositoryTemplate
  name: aux_cdb_templ
  container: PDB12C1
  reference: REPOSITORY_TEMPLATE_REF-4
  repository: OELC9/'/u01/app/oracle/12.1'
  template: aux_cdb_params
```

In the above example:

- a. **name** is what you want the template name to be.
 - b. the **container** is the source PDB.
 - c. the **repository** is the target environment/ORACLE_HOME that the CDB target is running from.
 - d. the **template** is the name of the database template created in the GUI.
3. Provision a virtual pluggable database (VPDB) via the Delphix GUI using the source PDB defined as the container in the repository template and provision the VPDB to the target environment defined by the repository (target environment/Oracle Home combination) in the repository template.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)
- [Configuration Settings for Oracle Virtual Databases](#)

Using the PurgeLogs Operation

This topic addresses scenarios in which the purgeLogs operation is used to clear expired and unused logs.

Log retention does not affect the logs in the following three cases:

- If the logs are expired but are required for snapshots to be consistent
- Logs generated after the most recent snapshot
- Logs required to make retention-proof bookmark points provisionable

It is possible that the storage space is not freed by retention. In this scenario, there is only one snapshot and under that circumstance, it would continue to generate logs to the point that there was no storage space available. This prevents new snapshots from being taken. Additionally, retention does not delete expired logs that are required for snapshots to be consistent.

 The purgeLogs operation does not function for CDBs or any non-Oracle data platforms at this time. In addition, this operation only functions for Oracle non-multitenant dSources, it does not function with VDBs.

Possible Strategies

Strategy 1: It deletes all the expired logs associated with snapshots as required until the target space to reclaim is achieved. It deletes the oldest snapshot logs and moves to the newest one in order for each Timeflow starting from the current Timeflow first. The logs required for retention proof bookmarks and pinned snapshots are not deleted.

Strategy 2: It deletes logs beyond the last snapshot to free up space in which retention cannot free. Deletion is done in this order as opposed to deleting from the oldest logs because of the possibility of Timeflow repair at a later time. The deletion is done based on the actual compressed file size of the log sequence on the Delphix Engine. If there is a retention proof bookmark, this operation preserves the logs required to provision from the retention proof bookmark. This operation also spares the online in-flight log sequence since chunks of this sequence will keep coming in (this only applies for logsync mode set to Archive + Online Redo).

PurgeLogs operation takes three arguments :

storageSpaceToReclaim - This takes a non-zero amount of space to reclaim (eg; 20K, 5M etc). This is the target amount of space that the user will want to free up. This will be dependent on whether there will be enough objects to be freed. By default, **Strategy 2** is always in effect. The user has the ability to run **strategy 1** if their target space reclaim is not met, the user can run **strategy 1** using the **deleteSnapshotLogs** argument.

deleteSnapshotLogs - This is set to false by default. If this is set to true, then logs are deleted based on **strategy 1**. Even after completing that, if the target storage space to reclaim is not achieved, then logs are deleted based on **strategy 2** until the target is achieved or there are no more logs to delete. If this is set to false, then deletion is done based on **strategy 2** only.

dryRun - This is a boolean argument and set to true by default. If it is true, then the operation does not actually delete the logs. This argument can be used to see which snapshots and parts of the Timeflow will be affected by the operation. This operation returns the time flow point beyond which logs might be deleted and the set of snapshots affected by this operation (and will not be provisionable afterwards). When the operation is run in non-dryrun mode, we compute the time range of the snapshot and provide the endpoint as the end of the truncated Timeflow. In dryrun mode, we return the endpoint of the first log which survived as the Timeflow point beyond which it has been truncated.

 • Use this operation with `setopt format=json`, otherwise, the list of snapshots may be truncated if there are more than 3.

- Step 7 is an optional step and to be used in the event **Strategy 2** does not clear enough space desired by the user. Please note that any snapshots affected by this step enabled can no longer be provisioned from. The user will be required to restore archivelogs to re-enable provisioning of the affected snapshot. To repair this function please use the TimeFlow Repair Tool.
- VDB transaction logs can be manually removed from the **archive** filesystem mounted on the target for each VDB.

This example shows how to free 9K of space from container "example_container" using strategies 1 and 2:

1. Log into the CLI
2. Go to the container using `/database`
3. Set the output format into json using this statement: `setopt format=json`
4. Select the container using `select example_container`
5. Select operation `purgeLogs`
6. Set the target space to reclaim to 9K using `set storageSpaceToReclaim=9K`
7. **(Optional Step)** `deleteSnapshotLogs=true` . **Warning:** this step should be used only in the event that Strategy 2 does not free enough space. This will remove the ability to provision from affected snapshots or require the restoration of archivelogs to enable provisioning. Use `set dryRun=true` to run operation and view the snapshots which will be affected. Affected snapshots will be listed under the section `affectedSnapshots` .
8. Use `set dryRun=false` to enable the operation.
9. `Commit` operation

The output should look like this:

```
delphixengine> /database
delphixengine database> select example_container
delphixengine database "example"> purgeLogs
delphixengine database "example" purgeLogs *> set storageSpaceToReclaim=9K
delphixengine database "example" purgeLogs *> set deleteSnapshotLogs=true
delphixengine database "example" purgeLogs *> set dryRun=false
delphixengine database "example" purgeLogs *> commit
  type: PurgeLogsResult
  affectedSnapshots:
    0:
      type: OracleSnapshot
      name: '@2014-08-05T23:29:57.576Z'
      consistency: INCONSISTENT
      container: dbdhcp3
      creationTime: 2014-08-05T23:29:57.576Z
      firstChangePoint:
        type: OracleTimeflowPoint
        location: 2825963
        timeflow: dbdhcp3/default
        timestamp: 2014-08-05T23:31:15.000Z
      latestChangePoint:
        type: OracleTimeflowPoint
        location: 2826079
        timeflow: dbdhcp3/default
```

```
timestamp: 2014-08-05T23:31:15.000Z
missingNonLoggedData: false
namespace: (unset)
reference: ORACLE_SNAPSHOT-75
retention: 0
runtime: (unset)
timeflow: dbdhcp3/default
timezone: America/New_York,EDT-0400
version: 11.2.0.2.0
truncatePoint:
  type: OracleTimeflowPoint
  location: 3176794
  timeflow: dbdhcp3/default
  timestamp: 2014-08-07T19:58:55.000Z
```

This output shows that the affected snapshot is '@2014-08-05T23:29:57.576Z' and the snapshot will not be provision-able. The truncate point indicates the last SCN/timestamp that can be provisioned.

Related Topics

- [Provisioning and Managing Virtual Databases with Oracle](#)

Advanced Linking

This section covers the following topics:

- [Enabling Validated Sync for Oracle](#)
- [Linking Oracle Physical Standby Databases](#)
- [Specifying External Data Directories for Oracle dSources and VDBs](#)
- [Linking to Oracle dSources with RMAN Compression or Encryption Enabled](#)
- [Upgrading dSources after an Oracle Upgrade](#)
- [Oracle Source Continuity](#)
- [Oracle LiveSources](#)
- [Oracle RAC dSource Node Addition or Deletion](#)
- [Using the Double Sync option for Oracle SnapSync](#)
- [Linking dSources from an Encrypted Oracle Database](#)
- [Detaching and Re-Attaching Oracle dSources](#)
- [Provisioning from a Replicated Oracle dSource](#)
- [Detaching and Re-attaching a PDB dSource from a Data Guard Site](#)
- [Working with Oracle Snapshots](#)
- [Moving the PDB to a New CDB](#)

Enabling Validated Sync for Oracle

This topic describes the validated sync process for Oracle databases using both the Delphix Management application and Command Line Interface (CLI).

Traditional Oracle dSource snapshots require some recovery during provisioning. By configuring validated sync for Oracle, the Delphix Engine selects a compatible Oracle installation and applies the recovery necessary to provision a snapshot immediately after each SnapSync. Snapshots that have been through this validated sync process step do not require recovery during provisioning.

 The Delphix Engine may be unable to perform validated sync on a physical standby database in Real-Time Apply mode. This is because the standby may apply changes before copying the logs that contain those changes. Without the logs necessary to perform recovery, validated sync cannot be executed. However, you can still provision the snapshot when the archive logs become available on the standby.

Prerequisite: Designating a Staging Host

In order to validate an Oracle dSource snapshot after a sync, the Delphix Engine requires a host with an Oracle installation that is compatible with the dSource. This machine is known as the **staging** host. You must explicitly designate which machines you want the Delphix Engine to use as staging hosts. All machines that have been marked as staging hosts are added to a pool. During sync validation, the Delphix Engine will select a compatible host from the pool, export the requisite archived redo logs and data files, and execute Oracle media recovery on the host. Follow these steps to designate a staging host.

1. Log into the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, select the environment you want to designate as staging.
5. Select the **Databases** tab.
6. Scroll down to the installations you want to designate as staging and edit the **Installation Details** by clicking on the **pencil** icon.
7. Select the **Use as staging** checkbox.
8. Select the green checkmark to confirm your change.

To configure validated sync for multiple dSources with different Oracle versions, you must designate a compatible staging source for each. If multiple compatible staging sites exist, the Delphix Engine will select one at random.

 The validated sync process will consume some resources on the staging host when snapshots are taken. Designating a performance-critical host as a staging host is not recommended.

 The default OS user for the staging host must have access to the Oracle installation that will be used to perform recovery during validated sync.

Oracle Pluggable Database

Validated sync for Oracle pluggable databases is not supported in this release.

Enabling Validated Sync

Oracle validated sync can be enabled at link time or on any existing dSource. When adding the dSource (at link time), in the Data Management tab, select Show Advanced, and select the Enable checkbox for Validated Sync.

For an existing dSource:

1. Log into the Delphix Management application using **Delphix Admin** credentials.
2. Click **Manage**.
3. Select **Datasets**.
4. In the **Datasets** panel, select the dSource for which you want to enable sync validation.
5. Select the **Configuration** tab and then select the **Data Management** sub-tab.
6. Select the pencil icon located next to the **VALIDATED SYNC CONFIGURATION**.
7. Select one of the following from the **Validated Sync Mode** drop-down list:
 - a. Transaction Log
 - b. Full or Differential
 - c. Full
 - d. None

The screenshot shows the configuration interface for SnapSync. The 'Configuration' tab is selected, and the 'Data Management' sub-tab is active. The 'Validated Sync' checkbox is checked. The 'LogSync' section shows 'Enabled' status and 'Archive + Online Redo' policy. The 'Nologging Diagnosis' section shows 'Nologging is being diagnosed'.

Shield [Pencil icon]

Timeflow Status **Configuration**

Source Policies **Data Management** Masking Hooks

SnapSync

GENERAL SETTINGS

Validated Sync

RMAN SETTINGS

Backup type

SCN-Based

Level

RMAN channels

5

Maximum files per backup set

2

Check logical corruption

NETWORK SETTINGS

Encrypted

Compressed

Bandwidth limit (MB/s)

0

Number of connections

1

[X] [Checkmark]

LogSync

Enabled

Policy

Archive + Online Redo

Nologging Diagnosis

Nologging is being diagnosed

8. Select the checkmark to confirm your change.

Related Topics

- [Advanced Linking](#)

Linking Oracle Physical Standby Databases

This topic describes special considerations for linking Oracle physical standby databases.

The Delphix Engine supports linking both physical and logical standby databases. In previous versions of the Delphix Engine, limitations were placed upon support for Oracle RAC physical standby databases in Real-Time Apply mode. In version 3.0 of the Delphix Engine, these restrictions were lifted.

Using Block Change Tracking (BCT) on a Physical Standby Database

In general, Delphix recommends enabling Block Change Tracking (BCT) on a primary or standby source database. See [Physical Standby Database Support Matrix](#) below for restrictions on enabling BCT on a standby database.

BCT is available from Oracle release 11.1.0.6 onward for physical standby databases **only** if they are licensed for the Active Data Guard option.

- Release 11.1.0.6 is unstable for the BCT on physical standby feature
- Release 11.1.0.7 requires a patch for Oracle bugs 7613481, 9068088
- Release 11.2.0.2 requires patches for Oracle bugs 10170431, 12312133
- Release 11.2.0.3 requires patches for Oracle bugs 12312133, 16052165,

Patches Required

Enabling BCT on a physical standby database without these patches is **not recommended** because of serious performance and stability issues.

BCT on a primary database has been stable since Oracle version 10.2.0.5. In order to make use of BCT (>11.2.0.4), The Physical Standby Database must be in a "Managed Recovery Mode", i.e. achieved using "ALTER DATABASE RECOVER **MANAGED** STANDBY DATABASE".

Physical Standby Database Support Matrix

Oracle Version	Apply Mode	Notes
10.2, 11.x and 12.x in Level Backup mode; 10.2 and 12.x in SCN Backup mode	Archive Apply mode	No special restrictions.
	Real-Time Apply mode	LogSync must be enabled.
11.x in SCN Backup mode	Archive Apply mode	<p>If the Physical Physical Standby Database is at version 11.2.0.4 or above, no special actions are required.</p> <p>Due to Oracle bug 10146187, Redo Apply must be stopped and the database opened in read-only mode during SnapSync. See the section <i>Stopping and Restarting Redo Apply</i> below for more information.</p>

Oracle Version	Apply Mode	Notes
	Real-Time Apply mode	<p>If the Physical Physical Standby Database is at version 11.2.0.4 or above, no special actions are required.</p> <p>LogSync must be enabled.</p> <p>Due to Oracle bug 10146187, Redo Apply must be stopped and the database opened in read-only mode during SnapSync. See the section <i>Stopping and Restarting Redo Apply</i> below for more information.</p> <p>In addition, to avoid Oracle Bug 13075226, which results in a hang during the restart of Redo Apply, Delphix requires disabling using BCT on the standby database. The hang occurs when BCT is enabled on a standby database that uses SCN backup mode.</p> <div data-bbox="842 875 1423 1541" style="border: 1px solid #f0e68c; padding: 10px;"> <p>⚠ Patch Required</p> <p>SnapSync will fail If running Oracle 11.2 before 11.2.0.4 when using SCN backups and real-time apply mode, Use level based backups instead.</p> <p>If the Oracle installation has already been patched for Oracle bug 13075226, or once the patch is applied, use the CLI to update the repository for this installation so that applied Patches include Oracle bug number 13075226. If the repository does not indicate that Oracle bug 13075226 for the repository has been addressed, SnapSync will not be possible when using SCN backups and real-time apply. See <i>Updating Repository for Oracle applied patches with the Command Line Interface</i> below for details on how to update the repository.</p> </div>
18.0.0.0 - 19.3.0.0	Real-Time Apply mode	Oracle patch 29056767 installed, Failure to install this patch may result in SnapSync failure with error.

Level Backup Mode for SnapSync

By default, the Delphix Engine's SnapSync feature uses **SCN Backup** mode and is designed to not interfere with other backups that may already be in use. However, in cases where RMAN is not being used outside of the Delphix Engine, the Delphix Engine can use the **Level Backup** mode that improves SnapSync behavior on Oracle 11g physical standby databases. In this mode, **redo apply** does not have to be stopped during SnapSync. See [Advanced Data Management Settings for Oracle dSources](#) for more information about SnapSync settings.

Requirements for Using Level Backup Mode

Customer not backing up their physical standby with RMAN:

- Set CONTROL_FILE_RECORD_KEEP_TIME to 365

OR all of the following:

- Physical standby database running Oracle 11.2.0.2 or later version
- All RMAN backups must use tags
- RMAN CROSSCHECK commands must specify tags
- RMAN DELETE commands must specify tags
- RMAN DUPLICATE commands must specify tags
- Set CONTROL_FILE_RECORD_KEEP_TIME to 365

 Failure to meet all of these requirements will cause external RMAN backups to be incomplete or result in corrupt SnapSync snapshots. Switching from SCN to LEVEL mode will force a new LEVEL 0 backup.

Stopping and Restarting Redo Apply

Oracle bug 10146187 requires stopping of redo apply before an SCN-based incremental backup can be issued. These scripts can be used as pre- and post-scripts during the dSource linking process to stop and restart **Redo Apply**.

- SnapSync pre-script: [disableStandby.sh.template](#)
- SnapSync post-script: [enableStandby.sh.template](#)

These scripts must be modified for local use, particularly in regard to whether the physical standby database operates in MOUNTED or OPEN mode.

 Failure to properly customize these scripts could violate your Oracle license terms by running redo apply on an open database, which requires an Oracle Active Data Guard license.

Linking and Provisioning a Mounted Standby

Warning

When you link a standby database in the mounted mode and have not been opened read-only, the data files for temporary tablespaces will be present in v\$tempfile but will not actually be created yet, and therefore will indicate 0 file size. As a result, any VDB provisioned from a snapshot taken in this state will end up with the same tablespaces and datafiles created, but with a default file size of 52428800 bytes.

For databases that are in the **mounted** state, the Delphix database user account must be **SYS** (having the **SYSDBA** role), **SYSBACKUP** (having the **SYSBACKUP** role) or **SYSDG** (having the **SYSDG** role).

 **SYSBACKUP** and **SYSDG** roles are only available in Oracle 12.1 and later releases.

However, for an **open** standby (Active Data Guard) database, only a regular database user account is required.

Connecting to a **mounted** standby with a **SYS** user account requires that the mounted standby be configured with a password file. Delphix does not capture the password file during SnapSync, and for this reason, cannot provision or sync validate a database with a SYS user. A secondary, regular database user account can be specified through either the **Delphix Management** application or CLI. This database user will then be used to connect to the database during provisioning and validated sync. Note that the **SYS** user is still required to perform snapshots of the source database.

In the **Delphix Management** application, the **non-SYS user** can be specified from within the **Add dSource** wizard, or on the back of the Oracle dSource after linking.

Configuring a Standby PDB in Mount Mode

 This procedure is optional and only applies if SCM is disabled.

To configure a standby PDB in the mount mode, you must also provide a non-SYS user for both the CDB and the PDB. The PDB non-SYS user can only be added via the CLI. You must perform a fresh SnapSync after adding the non-SYS user.

Pre-requisite

In order for Delphix Engine to connect, you must configure a static listener configuration for the PDB. You can configure a static listener by adding a configuration into listener.ora and restarting the listener.

```
SID_LIST_LISTENER=
(SID_DESC=
(GLOBAL_DBNAME=CDOMLOS4F71PDB1)
(SID_NAME=stby18c)
(ORACLE_HOME=/u01/app/oracle/product/18.0.0.0/dbhome_1)
)
)
```

In the above example configuration, GLOBAL_DBNAME is the PDB name and SID_NAME is the SID of the CDB.

Procedure

Run the following commands to configure a PDB and CDB in the mount mode.

1. Update PDB non-SYS user.

```
# Update PDB nonsys user
delphix> /sourceconfig
delphix sourceconfig> select RH74PDB04
delphix sourceconfig 'RH74PDB04'> update
delphix sourceconfig 'RH74PDB04' update *> set nonSysUser=delphix
delphix sourceconfig 'RH74PDB04' update *> set
nonSysCredentials.type>PasswordCredential
delphix sourceconfig 'RH74PDB04' update *> set
nonSysCredentials.password=delphix
delphix sourceconfig 'RH74PDB04' update *> commit;
```

2. Update CDB non-SYS user.

```
# Update CDB nonsys user
delphix> /sourceconfig
delphix sourceconfig> select rh74cdb2
delphix sourceconfig 'rh74cdb2'> update
```

```
delphix sourceconfig 'rh74cdb2' update *> set nonSysUser=delphix
delphix sourceconfig 'rh74cdb2' update *> set
nonSysCredentials.type>PasswordCredential
delphix sourceconfig 'rh74cdb2' update *> set
nonSysCredentials.password=delphix
delphix sourceconfig 'rh74cdb2' update *> commit;
delphix sourceconfig 'rh74cdb2'>
```

3. Perform sync of the PDB.

```
# Perform sync of PDB
delphix> /database
delphix database> select RH74PDB04
delphix database 'RH74PDB04'> sync
```

Setting the Non-Sys User on the Oracle dSource

1. Create the delphix_db user in the primary database.
2. Log into the Delphix Management application.
3. From the **Manage** menu, select **Datasets**.
4. From the Configuration tab select the Oracle dSource for which you want to add a **non-SYS** user.
5. Click the dSource's icon to open the dSource information pane.
6. Click the Edit button next to **Non-SYS User**.
7. Enter a non-SYS user and credentials that exist on the standby.
8. Click the **Accept** button to save this user and associated credentials.

The non-SYS user will be used to connect to all VDBs provisioned from snapshots of this dSource that are created after the non-Sys user has been set.

Updating repository for applied patches with the Command Line Interface

1. Select the repository of the database

```
delphix> repository select '/opt/app/oracle/product/11.2.0.2/db_1'
```

2. Execute the `update` command.

```
delphix repository '/opt/app/oracle/product/11.2.0.2/db_1'> update
```

3. Set `appliedPatches` to list current patches applied to the repository.

```
delphix repository '/opt/app/oracle/product/11.2.0.2/db_1'update *> set
appliedPatches=13075226
```

4. `Commit` the operation.

```
delphix repository '/opt/app/oracle/product/11.2.0.2/db_1'update *> commit
```

Setting the Non-Sys User with the Command Line Interface

1. Select the `source config` of the mounted standby.

```
delphix> sourceconfig select pomme
```

2. Execute the `update` command.

```
delphix sourceconfig "pomme"> update
```

3. Set the `nonSysUser` and `nonSysCredentials` to a non-SYS user that exists on standby.

```
delphix sourceconfig "pomme" update *> set nonSysUser=<non-sys-username>
delphix sourceconfig "pomme" update *> set
nonSysCredentials.type>PasswordCredential
delphix sourceconfig "pomme" update *> set nonSysCredentials.password=<non-sys-
password>
```

4. Commit the operation.

```
delphix sourceconfig "pomme" update *> commit
```

Related Topics

- [Data Management Settings for Oracle Data Sources](#)
- [Advanced Linking](#)

Specifying External Data Directories for Oracle dSources and VDBs

This topic describes the process for including external data files with dSource snapshots and VDBs.

In the following places, you can specify the directory for any external data files that should be included with dSource snapshots:

- During the dSource linking process click on the **Advanced** section of the **Data Management** screen
- After you have created the dSource go to the **Configuration** tab

External File Import for the Delphix Engine and VDBs

The Delphix Engine will not fetch external tables or external data types such as BFILE. Instead, in order to link external data files to the source database and make it available to the Delphix Engine, you must create a directory in the file system and the database. Any data files in the directory you specify will be applied, recursively, to the dSource.

External data will be provisioned to each VDB that is created from this dSource. You will need to update the external file/data type definition to point to the new location after creating VDBs. Provisioning a VDB with external data creates a directory named **external** in the VDB mount point location.

Configuring the rsync Command Location for an Environment

Files from the external data directory are fetched using the rsync command installed in the source environment. In order to SnapSync a dSource with an external data directory, rsync must be installed in the source environment. If rsync is installed in a non-standard location, the path to the rsync command can be configured in the **Environment Details** for the source environment on the **Environment Management** screen.

Example of Attaching and Redirecting External Data Files for Oracle Databases

This example uses two environments:

1. **172.16.200.446** as the source environment
dinosaur as the source database
2. **172.16.200.447** as the target environment
vdino as the target database

Linking a dSource

1. Create an external data directory and an external data file, and attach the directory to the source database.

- a. Log into **172.16.200.446** as the environment user.
- b. Create a physical directory on the source environment.

```
$ mkdir /work/extdata
```

- c. Create a directory in Oracle.

```
$ sqlplus / as sysdba
SQL> create or replace directory extdata as '/work/extdata';
```

- d. Create a text file `/work/extdata/exttab.dat`.

```
$ cat > /work/extdata/exttab.dat
1, aaa
2, bbb
3, ccc
^C
```

- e. Create an external table `exttab` .

```
$ sqlplus / as sysdba
SQL> create table exttab (id number, text varchar2(10))
2 organization external (default directory extdata location('exttab
.dat'));
```

- f. Query the table.

```
SQL> select * from exttab;
   ID TEXT
-----
    1 aaa
    2 bbb
    3 ccc
```

2. During the process of linking the dSource to the **Dinosaur** database, or in the dSource's **Configuration** tab after creating the link, enter `/work/extdata` in the **External Data Directory** field.

Provisioning a VDB

1. Provision **vdino** from **Dinosaur**.
2. Modify the directory **extdata** in **vdino**
 - a. Log into the target environment **172.16.200.447**
 - b. Set **SID** to `vdino`

```
$ export ORACLE_SID=vdino
```

- c. A query to `exttab` will fail.

```
$ sqlplus / as sysdba
SQL> select * from exttab
select * from exttab
*
ERROR at line 1:
ORA-29913: error in executing ODCIEXTTABLEOPEN callout
ORA-29400: data cartridge error
KUP-04063: unable to open log file EXTTAB_23394.log
OS error No such file or directory
ORA-06512: at "SYS.ORACLE_LOADER", line 19
```

3. Modify directory to the new location.

```
SQL> create or replace directory extdata as '/mnt/provision/vdino/external';
```

4. Query `exttab` again.

```
SQL> select * from exttab;
  ID TEXT
-----
   1 aaa
   2 bbb
   3 ccc
```

Related Topics

- [Advanced Linking](#)

Linking to Oracle dSources with RMAN Compression or Encryption Enabled

This topic describes the behavior of the Delphix Engine when linking to a dSource with RMAN compression or encryption enabled.

In earlier versions of the Delphix Engine, the dSource linking process would fail if RMAN compression or encryption was enabled. In order for the linking process to complete, the administrator was required to ensure that compression was not enabled for device type **SBT_TAPE**, and that encryption was also not enabled.

With the Delphix Engine, linking a dSource succeeds if compression or encryption is enabled, but the RMAN backup that creates the dSource will not be compressed or encrypted. This is true in the case where the administrator has enabled compression for tape, and in the case where the administrator is using OSB and has enabled encryption for tape.

You can check the RMAN compression and encryption settings with the commands `show device type` and `show encryption for database`, respectively.

If compression is enabled on a source environment, then it may need to be enabled for the target environment. Compression should be enabled for target environments if all segments in a source environment utilize **NOCOMPRESS** or **COMPRESS="BASIC"**.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)

Upgrading dSources after an Oracle Upgrade

This topic describes how to upgrade dSources after an Oracle database upgrade.

Prerequisites

Do not suspend LogSync on the Delphix Engine during an Oracle upgrade of the source environment. LogSync will detect the Oracle version change. It will then notify you to refresh the host and use the Update icon on the configuration tab for all the associated dSources and VDBs (see below). Follow all Oracle instructions and documentation.


⚠ In order to upgrade your Oracle dSource from Non-Multi-Tenant (Non-MT) to Multi-Tenant(MT) database, please refer to [Prepare and Upgrade a Non-MT Oracle dSource to MT](#) section.

There are 2 ways to apply a PSU (Patch Set Update)/Oracle upgrade:

- Apply to existing ORACLE_HOME. Best if on Delphix v4.1.x or higher.
- Create a new ORACLE_HOME (or clone an existing one) and then apply the PSU to the new ORACLE_HOME

Applying to an existing ORACLE_HOME

1. Following Oracle documentation, patch the ORACLE_HOME and the database.
2. Click the refresh icon on the top right of the selected environment.



3. On the Refresh Environment pop-up window, click **Refresh**.
4. Under the **configuration** tab, go to the PDB and verify that the **Repository** and/or **Version** has been updated.

Creating a new ORACLE_HOME

1. Refresh the environment from the Delphix Management application.
2. In the **Databases** tab, verify that the new ORACLE_HOME appears as an ORACLE Installation.
3. Following Oracle documentation, patch the production database, etc.
4. Click **Manage**.
5. Select **Datasets**.
6. Expand the group(s) containing all the non-multitenant and multitenant dSources.
7. Select **dSource**.
For upgrading a PDB or multitenant dSource, select the container CDB **dSource**.
8. Click the **Configuration** tab.
9. From the **Actions** menu (...), select **Upgrade** to switch the ORACLE_INSTALLATION to the new one.
10. Under the **configuration** tab, go to the PDB and verify that the **Repository** and/or **Version** has been updated.

✓ Updating the Oracle User after an Upgrade

There may be cases when you upgrade the Oracle home and the Oracle User (the OS user who owns the Oracle binary) is a different user than the previous Oracle User. You will then need to update the Oracle User for each environment, and then re-connect each dSource and VDB to the upgraded Oracle home

using the new Oracle User.

The new Oracle User must be in the same OS group (for example, dba or oinstall) as the previous one.

1. Log into the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **environment** where you want to add the user.
5. Next to **Environment Users**, click the **Pencil** icon to edit the new user.
6. Follow the procedure to upgrade the dSources and VDBs described in this topic.

Post-requisites

After upgrading the dSource to a new major release of Oracle (11.2.0.4 to 12.1 for example), you must re-run the [createDelphixDBUser.sh](#) script.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)

Oracle Source Continuity

Overview

In earlier versions of the Delphix Engine, when an Oracle database underwent a resetlogs operation, you were required to re-link the Oracle source. This meant that you had to completely back up the Oracle database and store it again on the Delphix Engine. If any virtual databases (VDBs) were provisioned from the dSource and needed to be saved, you had to rename and save the old dSource, resulting in a possible doubling of storage space consumed on the Delphix Engine. The old VDBs could not be refreshed to the relinked dSource.

Beginning with Delphix Engine version 4.1.1.0/4.0.6.0, the Oracle database no longer requires you to re-link sources after a resetlogs operation. The Delphix Engine will detect this condition, automatically take a new full backup, and create a new TimeFlow for the next SnapSync of the source. Benefits of the Oracle Source Continuity feature include:

- Lower storage costs and easier administration.
 - Only the changed blocks of the new SnapSync backup will be stored on the Delphix Engine. Because of the way the Delphix Engine handles duplicate blocks, the full backup likely has a storage requirement similar to an incremental backup.
- Existing VDBs provisioned from previous snapshots for the source will remain.
 - You can use and refresh those VDBs to the new snapshot

The improved user workflow replaces the old user workflow, which directed users to troubleshoot when SnapSync would fail. Begin Oracle Source Continuity in the following way:

1. The database undergoes a resetlogs operation.
2. If LogSync is enabled, it generates a fault and stops.
3. Start SnapSync. The SnapSync does a full restore of the database to a new TimeFlow, clears the fault, and restarts LogSync. If you created VDBs prior to the resetlogs operation, they will still exist after the SnapSync; you can refresh them from the new snapshot.

Creating a New TimeFlow

When LogSync detects the resetlogs operation, it will still stop and generate a fault. LogSync must stop because a new timeline has been created on the database. This usually happens because the database has been rewound to a past point. The transaction logs being generated on the new timeline are out of sync and conflict with logs from the old timeline. The data files are also out of sync with the data files on the Delphix Engine. You must create a corresponding new Timeflow on the Delphix Engine to store the new logs and new versions of the data files. This requires taking a new backup of the database.

Once LogSync detects the reset operation and throws the fault, no more changes will be retrieved from the database until you start a new SnapSync. This SnapSync will take a full backup, clear the fault, and restart LogSync. Only the new snapshot and Timeflow will be visible in the dSources view. Previous snapshots and Timeflow will still exist and be visible through the command-line (CLI) Capacity screen.

The following CLI output shows that the old and new Timeflow and snapshots are still available. The name of the original Timeflow for the database is "default." The name of the new Timeflow that was created during the SnapSync is "CLONE@2015-01-15T17:07:20."

```
delphix> /TimeFlow list display=name,container
```

NAME	CONTAINER
'CLONE@2015-01-15T17:07:20'	dbdhcp1
default	dbdhcp1

```
delphix> /snapshot list display=name,container,TimeFlow
```

NAME	CONTAINER	TimeFlow
'@2015-01-16T00:50:08.784Z'	dbdhcp1	default
'@2015-01-16T00:52:13.685Z'	dbdhcp1	default
'@2015-01-16T00:53:46.873Z'	dbdhcp1	default
'@2015-01-16T00:55:18.079Z'	dbdhcp1	default
'@2015-01-16T01:08:02.411Z'	dbdhcp1	'CLONE@2015-01-15T17:07:20'

The old snapshots and Timeflow will still be subject to logfile and snapshot retention policies. You can also delete the snapshots manually. In addition, you can use the CLI to provision from the old TimeFlow.

Related Topics

- [Managing Oracle Environments and Hosts](#)

Oracle LiveSources

Oracle LiveSources Overview

Prior to Delphix Engine version 4.2, users ran reports against virtual databases (VDBs) that they created with the Delphix Engine. Although this workflow helped them offload the reporting load from production, the data in the VDBs was not updated asynchronously. If users wanted newer data, they had to stop their reporting applications, refresh their VDBs, and resume. In the current release, you can run reports against data that is constantly being updated. There is one live data feed for each source database that is linked as a dSource on the Delphix Engine. You can point your reporting applications to this live feed. Additionally, you will continue to have all existing Delphix functionality from the dSource, such as creating read/write VDBs.

Understanding Oracle LiveSources

Oracle LiveSources leverage native Oracle Active Data Guard technology to keep a standby database up-to-date with changes happening on the source. The standby database is kept open for reads while it applies changes from the source. You can now connect to this standby database for real-time reporting needs. Using Delphix in conjunction with Active Data Guard gives you the ability to get both live up-to-date data and historical points in time from which you can provision virtual databases.

Understanding How to Use Oracle LiveSources

Oracle LiveSources Provide a Read-Only Live Data Stream from Delphix

You can convert an Oracle dSource to a LiveSource, which is a real-time read-only feed of the linked source. You can access the LiveSource using a JDBC string. Internally, a LiveSource is a standby database instance tracking the Linked Source in real-time managed mode and opened in read-only mode.

Understanding Oracle LiveSources with Data Age and Threshold

One of the important utilities of a LiveSource is that it provides a real-time feed of the linked source. In some instances, due to slow networks or other reasons, the LiveSource might fall behind the linked source that it is tracking. When adding a LiveSource, you can specify a data age threshold. If the LiveSource falls behind the linked source by more than the data age threshold, the Delphix Engine will generate a fault and inform you.

The LiveSource information is displayed on the Status tab. Delphix continuously monitors the standby instance and notifies you of any abnormalities.

- Live Data Status - Indicates that the LiveSource standby database instance is actively receiving data from the source database
- Data Age - Displays the data age of the LiveSource

You can change the Data Age Threshold at any time by updating the threshold value located on the **Configuration > Data Management** tab.

Oracle LiveSources Quickly Sync with Consistent Snapshots

Taking snapshots of a LiveSource is instantaneous since the standby database for the LiveSource is constantly receiving data from the source database and recovering it. Taking snapshots occurs instantaneously by taking a filesystem level snapshot of the data on the Delphix Engine without requiring an RMAN backup of the source database. All LiveSource snapshots are consistent; as a result, provisioning from LiveSource snapshots is fast, because no database recovery needs to happen.

Oracle LiveSources Use Resync and Apply

Resync is a way to refresh the LiveSource to the current point in the linked source. The following situations require a Resync to be performed:

- There are unresolvable gaps in the log sequence – for example, logs from the source database deleted before the primary database could ship them over to the LiveSource standby.
- The source database was taken through a point in time recovery/flashback, resulting in a changed incarnation.
- The source database contains non-logged changes. In this case, a Resync is needed only if you are interested in moving the non-logged data over to the LiveSource.
- The LiveSource is significantly behind the source database due to network communication issues or a large number of writes.

LiveSource Resync is a two-step operation of:

- **Start LiveSource Resync** – Start Resync performs an incremental backup of the source database to transfer the latest changes to the Delphix Engine. This operation does not affect the availability of the LiveSource
- **Apply LiveSource Resync** – Applying the Resync data will perform one more incremental backups from the source database to ensure up-to-date data, and recreate the LiveSource instance while preserving all the configurations. This operation requires downtime for the LiveSource.
- **Discard LiveSource Resync** – If the prepared resync data is no longer needed or resync has become obsolete (for example, another controlled change has been done on the source database), you can discard the current resync data with Discard LiveSource Resync. The next Resync will refetch data from the source database.

Pre-requisites: Configuration and Installation of Staging Environments To Host a Standby Database

Oracle Active Data Guard Required

The LiveSource feature requires an Active Data Guard license. Delphix uses Active Data Guard to replicate changes from the source database that it creates on the staging environment.

Network Requirements

LiveSource requires a Data Guard connection between the source and the standby database which utilizes TNS listeners associated with the databases.

Database Requirements

LiveSource requires Enterprise Edition of Oracle Database.

Oracle Support Matrix	
LiveSource Supports:	LiveSource Does Not Support:
Oracle 11g and non-multi tenant Oracle 12c, 18c, and 19c	Oracle Standard Edition on the source and staging environments
Physical and standby source databases	LiveSources running on a RAC

Related Topics

- [Oracle LiveSource User Workflows](#)


Oracle LiveSource User Workflows

Please use the following documentation as a guide to identify and act on common Oracle LiveSource User Workflows. The following table of contents includes steps for how to convert a dSource into a LiveSource, a provision from a LiveSource, sync a LiveSource, convert a LiveSource back to a dSource, and many other data procedures.

Converting to LiveSource from a dSource

To get a live feed to the source database data through the Delphix Engine, you must first link the database to the Delphix Engine to create a dSource. You can then convert the dSource into a LiveSource by following the steps outlined below:

1. In the left-hand panel, click the **dSource**.
2. From the Actions menu (...) select **Convert to LiveSource**, as highlighted below. This launches the **Convert to LiveSource** wizard.



Convert to LiveSource Database wizard

1. In the Overview tab select **Next**.
2. Enter a **DB Unique Name** for the LiveSource.
3. Enter a **Database SID** for the LiveSource.
4. Click **Next**.

i The LiveSource database name must be the same as the database name of the primary database; therefore, this value is read-only.

Convert to LiveSource, Environment tab

Select the environment on which the LiveSource will be created:

1. Select an **Environment User** for the LiveSource instance.
2. Enter the **Mount Point** for the LiveSource instance.
3. Select **Listeners** as needed. If you enable **Auto Select Listeners**, the Delphix Engine will pick the first available listener from the environment.

4. Click **Next**.

Convert to LiveSource, Configuration Template tab

1. Select **VDB configuration templates** for the LiveSource.
2. Enter additional **DB configuration parameters** for the LiveSource.
3. Click **Next**.

Convert to LiveSource, Data Management tab

1. Enter the **data age warning threshold** for the LiveSource.
If the data in LiveSource lags behind the source database by more than this threshold, the Delphix Engine will raise a fault and notify you.
2. Click **Next**.

Convert to LiveSource, Hooks tab

1. Enter the **operations** to be performed on initial conversion. These operations are performed after the Delphix Engine has created the standby database for the LiveSource.
2. Click **Next**.

 These operations will also be performed when resyncing a LiveSource.

Convert to LiveSource, Summary tab

1. Review the configuration summary.
2. Click **Convert** to begin the conversion.

Setting up Log Transport between a dSource or Primary Database and a LiveSource or Standby Database

After adding a LiveSource instance, you must configure the log transport between the dSource or primary database and the LiveSource or standby database. For details on configuring a standby database, refer to the Oracle Data Guard Concepts and Administration guide.

At source/primary database

1. Configure the `LOG_ARCHIVE_CONFIG` parameter to enable the sending of redo logs to remote destinations and the receipt of remote redo logs (the LiveSource instance). For example:


```
alter system set log_archive_config='DG_CONFIG=(sourcedb,livesource)'
scope=both;
```
2. Configure the `LOG_ARCHIVE_DEST_n` parameter to point the redo logs to the LiveSource instance. For example:


```
alter system set
log_archive_dest_2='SERVICE="(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=
livesource.dcenter.delphix.com)(PORT=1521))
(CONNECT_DATA=(SERVICE_NAME=livesource)(SERVER=DEDICATED)))" ASYNC
VALID_FOR=(ONLINE_LOGFILE,PRIMARY_ROLE) DB_UNIQUE_NAME=livesource'
scope=both;
```
3. Create a **passwd file** for the LiveSource into the target site.

4. Configure the corresponding `LOG_ARCHIVE_DEST_STATE_n` parameter to identify whether the log transport is enabled. For example:

```
alter system set log_archive_dest_state_2='ENABLE' scope=both;
```

5. Configure the `STANDBY_FILE_MANAGEMENT` parameter to enable automatic standby file management. For example:

```
alter system set standby_file_management='AUTO' scope=both;
```

RAC Instance as Standalone Environment

For RAC instances that are discovered as Standalone type environments, please apply the following steps instead of the steps above.

1.

```
alter system set log_archive_config='DG_CONFIG=(sourcedb,livesource)' sid='*' scope=both;
```
2.

```
alter system set log_archive_dest_2='SERVICE=livesource ASYNC VALID_FOR=(ONLINE_LOGFILE,PRIMARY_ROLE) DB_UNIQUE_NAME=livesource' sid='*' scope=both;
```
3. Set up **tnsnames.ora** in both source and target sites.
4. Make sure to have the same password file in all RAC instances and target site.
5.

```
alter system set log_archive_dest_state_2='ENABLE' sid='*' scope=both;
```
6.

```
alter system set standby_file_management='AUTO' sid='*' scope=both;
```

i The Delphix DB user must be updated to include grants for select on `v$managed_standby` and `v$dataguard_stats` in order to access MRP information and DataGuard statistics:

```
grant select on v_$managed_standby to <delphix DB username>;
create synonym <delphix DB username>.$v$managed_standby for v_$managed_standby;

grant select on v_$dataguard_stats to <delphix DB username>;
create synonym <delphix DB username>.$v$dataguard_stats for v_$dataguard_stats;
```

At the Staging Environment where the LiveSource standby database environment is running

1. Configure the `FAL_SERVER` parameter to point to the primary database for proper fetch archive log function. For example:

```
ALTER system SET fal_server='service="(DESCRIPTION=(ADDRESS=(PROTOCOL=tcp) (HOST=sourcedb.dcenter.delphix.com) (PORT=1521)) (CONNECT_DATA=(SERVICE_NAME=sourcedb) (SERVER=DEDICATED)))"' ;
```

2. If not already created, configure a password for Data Guard.

Removing a LiveSource

1. In the **Datasets** panel, click the **LiveSource**.
2. Click the **Configuration** tab.
3. From the Actions menu (...) select **Convert to dSource**.
4. Click **Convert to dSource**.

Taking a Snapshot on a LiveSource

To take a snapshot of a LiveSource:

1. In the **Datasets** panel, select the **LiveSource**.
2. In the upper right-hand corner, click the **camera** icon.

LiveSource snapshots are instantaneous, Quick Provision snapshots. They do not require an RMAN backup of the source database

Provisioning from a LiveSource TimeFlow

Provisioning from a LiveSource Timeflow is the same process as provisioning from a snapshot for dSource Timeflow. The only difference is that you will select a LiveSource and a LiveSource snapshot.

Enabling, Disabling, and Detaching a LiveSource

A LiveSource is **enabled** the same way as a regular dSource.

1. Login to the Delphix Management application as **admin** or another user with administrative privileges.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **LiveSource** you want to enable.
5. From the Actions menu (...) select **Enable**.
6. Click **Enable** to confirm.

i When you enable the LiveSource, the Delphix Engine will recreate the standby database on the staging environment.

A LiveSource is **disabled** the same way as a regular dSource. Disabling a LiveSource will stop further operations on the Delphix Engine related to the LiveSource.

1. Login to the Delphix Management application as **admin** or another user with administrative privileges.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **LiveSource** you want to disable.
5. In the upper right-hand corner, from the **Actions** menu (...) select **Disable**.
6. In the Disable dialog select **Disable**.

When you are ready to enable the LiveSource again, from the Actions menu (...) select **Enable**, and the LiveSource will continue to function as it did previously.

i Disabling a LiveSource shuts down the standby database that Delphix manages on the staging environment.

You can detach a LiveSource in the same way as [detaching a regular dSource](#). Detaching a LiveSource will implicitly convert the LiveSource into a regular dSource. After a dSource is re-attached, you can convert it back to a LiveSource.

Resyncing a LiveSource and Applying the Resync

Resync is a way to refresh the LiveSource to the current point in the linked source. Resync is a multi-phase operation comprised of the following:

Perform Resync

1. Click **Manage**.
2. Select **Datasets**.
3. Select your LiveSource.

4. From the Actions menu (...) select **Start LiveSource Resync**. The LiveSource can stay up while the Resync is in progress.

Discarding Resync Data

Prerequisites

- Resync is started and ready to apply

After Resync has finished, you can choose to not apply but rather discard the data that was brought over from the source database as part of Resync.

Procedure

To discard the data:

1. Click **Manage**.
2. Select **Datasets**.
3. Select your LiveSource.
4. From the Actions menu (...) select **Discard LiveSource Resync**.

Applying Resync Data

Prerequisites

- Resync started and ready to apply

Procedure

1. Click **Manage**.
2. Select **Datasets**.
3. Select your LiveSource..
4. From the Actions menu (...) select **Apply LiveSource Resync**.
5. If the apply resync data process fails, first investigate and resolve the cause of failure, such as a full disk. Then follow the procedure to start resync.

Migrating a LiveSource

1. Click **Manage**.
2. Select **Datasets**.
3. Select your LiveSource.
4. From the Actions menu (...) select **Disable**.
5. From the Actions menu (...) select **Migrate**.
6. Update the environment, user, and repository. From the **Configuration** tab select the **Source** sub-tab and click the Pencil (edit) icon next to the Database.
7. Enable the **dSource**.

Note: After the LiveSource is migrated to a different staging environment, you must ensure that the log transport between the source database and the LiveSource instance on the new staging environment is set up correctly.

Related Topics

- [Oracle LiveSources](#)

Oracle RAC dSource Node Addition or Deletion

Managing Oracle RAC Node Changes in Delphix

The addition and removal of cluster nodes from an Oracle Cluster must be considered where Delphix has dSources attached to RAC databases running from the same cluster environment or RAC VDB's provisioned out to those cluster nodes. Each scenario needs to be considered in its own right as dSources and VDBs require specific changes be made to accommodate the removal of a cluster node.

The changes required by Delphix and made within Oracle that need to be considered when the Oracle cluster node configuration alters and the process for removing or adding an Oracle cluster node is best demonstrated through an example.

The oracle cluster is comprised of 2 cluster nodes as seen from olsnodes.

```
[oracle@oelc9n1 ~]$ su - grid
```

```
Password:
```

```
Last login: Thu Nov 30 17:15:36 AEDT 2017 on pts/1
```

```
[grid@oelc9n1 ~]$ . oraenv
```

```
ORACLE_SID = [grid] ? +ASM1
```

```
The Oracle base has been set to /u01/app/grid
```

```
[grid@oelc9n1 ~]$ olsnodes -s -t
```

```
oelc9n1    Active    Unpinned
```

```
oelc9n2    Active    Unpinned
```

Node 2 in the cluster oelc9n2 is the node targeted for removal from the Oracle cluster.

Delphix dSources and Removing Oracle Cluster Nodes

In this example, Delphix has a 2 Node Oracle RAC database being ingested as a dSource. The sample database name is db121 comprised of 2 RAC instances db1211 and db1212. The instance impacted by the removal of node oelc1n2 is db1212. The dSource configuration appears as follows:

The screenshot displays the configuration page for dataset **db121**. The left sidebar shows a tree view with **db121** selected under **dSources**. The main content area is divided into several sections:

- Database:**
 - Unique Name: db121
 - Name: db121
 - Version: Oracle 12.2.0.1.0
 - Size: 1.36GB
 - User: delphix
- Environment:**
 - Name: oelc9n
 - OS: Linux (CentOS)
 - Timezone: America/New_York,EST-0500
 - User: oracle
 - Repository: /u01/app/oracle/product/12.2.0.1/dbhome_1
- Cluster Instances:**

Node Name	Instance N...	Instance N...
oelc9n1	1	db1211
oelc9n2	2	db1212

The cluster node, its associated listeners and instances must be deleted from the cluster through Oracles RAC RDBMS and Clusterware software using the processes described in the Oracle documentation.

This begins with the removal of the RAC instance using the process detailed in the Oracle RAC documentation. This example is operating against an Oracle 12.1 RAC and Cluster and the documentation link this specific release is the following:

Removing RAC Instances and Homes

<https://docs.oracle.com/database/121/RACAD/GUID-CE995C7B-B420-4A4E-B4EA-E7A972277588.htm#RACAD7903>

Removal of the RAC instance is typically executed through the database configuration assistant dbca and using the UI provided through dbca the instance db1212 is removed from node oelc9n2. After the instance has been deleted using dbca, Oracles srvctl utility will show one instance only (db1211) running from node oelc9n1.

```
[oracle@oelc9n1 ~]$ srvctl config database -d db121
```

```
Database unique name: db121
Database name: db121
Oracle home: /u01/app/oracle/12.1
Oracle user: oracle
Spfile: +DATA/DB121/PARAMETERFILE/spfile.289.961441955
Password file: +DATA/DB121/PASSWORD/pwddb121.277.961441737
Domain:
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: FRA,DATA
Mount point paths:
Services:
```

Type: RAC
 Start concurrency:
 Stop concurrency:
 OSDBA group: dba
 OSOPER group:
 Database instances: db1211
 Configured nodes: oelc9n1
 Database is administrator managed

Following the removal of the instance, the inventory on the node to be removed is updated to reflect the removal of the Oracle Home that was running this dSource instance on the node. This too must be performed according to the Oracle documentation. Oracles utility runInstaller is used for Oracle inventory maintenance. This must be performed using the operating system owner of the RDBMS Home (in this case oracle)

With the RAC instance and RAC Home removed the node itself can be removed using the process detailed in Oracles Clusterware Administration documentation,


Removing Oracle Cluster Nodes

<https://docs.oracle.com/database/121/CWADD/GUID-8ADA9667-EC27-4EF9-9F34-C8F65A757F2A.htm#CWADD90992>

This must be performed as the Oracle Clusterware home operating system owner (in this case grid).

Removal of the node from within Oracles RDBMS and Clusterware software does not alter the fact that Delphix still sees the Oracle Cluster and RAC configuration as it was the last time an environment refresh was performed against the cluster.

The Environment in Delphix shows both cluster nodes are still present along with any node listeners that may have been discovered running on that node.



The dSource shows the node and both instances still present this needs to be cleaned up from a Delphix perspective.

The screenshot displays the Delphix Management application interface. The top navigation bar includes 'Manage', 'Resources', 'System', and 'Help'. The left sidebar shows 'Datasets' and 'dSources' sections. The main content area is titled 'db121' and shows the 'Configuration' tab. The 'Source' sub-tab is active, displaying the following details:

- Database:** Unique Name: db121, Name: db121, Version: Oracle 12.2.0.1.0, Size: 1.36GB, User: delphix
- Environment:** Name: oelc9n, OS: Linux (CentOS), Timezone: America/New_York,EST-0500, User: oracle, Repository: /u01/app/oracle/product/12.2.0.1/dbhome_1


The 'Cluster Instances' table is shown below:

Node Name	Instance ...	Instance Name
oelc9n1	1	db1211
oelc9n2	2	db1212


Aligning Delphix with the New Oracle Cluster Configuration

In order for Delphix to pick up on the latest changes to the environment (OELC9) in this case, the removal of an entire node from the Oracle Cluster an environment refresh must be performed.

This is performed through the environments management panel in the Delphix Management application and the refresh icon present on the top right-hand side of the panel (highlighted in red).



Clicking on the icon presents a prompt indicating that an environment refresh is about to be initiated.



Click on refresh to kick off the environment refresh action.

The environment refresh will force Delphix to re-examine the Oracle Cluster.

This will result in the environment and the dSource being updated based on the cluster's current node, instance and listener configuration where only one node oelc9n1 is present.

After the node removal, the cluster environment in Delphix will look as follows and only oelc9n1 appears in the cluster node list.

The screenshot displays the Delphix Management interface for environment 'oelc9n'. The left sidebar shows the 'Environments' list with 'oelc9n' selected. The main content area is divided into several sections:

- Basic Information:** Cluster Name: oelc9n; Cluster Home: /u01/app/12.2.0.1/grid; Version: 12.2.0.1.0; Cluster User: oracle.
- Environment Users:** A table with columns 'Primary' and 'Name'. The 'Name' column contains 'oracle' with a star icon in the 'Primary' column.
- Faults:** A large box containing the text 'No Faults'.
- Cluster Nodes:** A list containing 'oelc9n1', which is highlighted in blue.
- Attributes:** Host Address: oelc9n1.dcenter; NFS Addresses: (empty).

The dSource will now be comprised of only one instance db1211 running from node oelc9n1.

The screenshot displays the Delphix Management application interface. The top navigation bar includes 'DELPHIX MANAGEMENT', 'Manage', 'Resources', 'System', and 'Help'. The main content area is divided into a left sidebar and a main panel. The sidebar shows a 'Datasets' section with a filter set to 'none' and a list of 'dSources' including 'db121' (Active) and 'VDBs'. The main panel shows the configuration for 'db121' under the 'Configuration' tab. The configuration is split into two columns: 'Database' and 'Environment'. The 'Database' column lists: Unique Name (db121), Name (db121), Version (Oracle 12.2.0.1.0), Size (1.36GB), and User (delphix). The 'Environment' column lists: Name (oelc9n), OS (Linux (CentOS)), Timezone (America/New_York,EST-0500), User (oracle), and Repository (/u01/app/oracle/product/12.2.0.1/dbhome_1). Below the configuration is a 'Cluster Instances' table.

Node Name	Instance N...	Instance N...
oelc9n1	1	db1211

Adding Oracle Cluster Nodes

You can add a node through the **Environments** management panel in the Delphix Management application and the **Refresh** icon present on the top right-hand side of the panel. Clicking on the icon presents a prompt indicating that an environment refresh is about to be initiated.

DELPHIX MANAGEMENT Manage Resources System Help Actions 0

Environments Filter: none

oelc9n Oracle Cluster - Enabled

Details Databases

Basic Information

Cluster Name
oelc9n

Cluster Home
/u01/app/12.2.0.1/grid

Version
12.2.0.1.0

Cluster User
oracle

Environment Users

Primary	Name
★	oracle

Faults
No Faults

DELPHIX MANAGEMENT Manage Resources System Help Actions 0

Environments Filter: none

oelc9n Oracle Cluster - Enabled

Refresh Environment

Refreshing an environment will rediscover all configured databases. To do so, enabled dSources may be temporarily disabled during the refresh. Are you sure you want to refresh this environment?

Cancel Refresh

Version
12.2.0.1.0

Cluster User
oracle

Environment Users

Primary	Name
★	oracle

No Faults

Click on the **Refresh** option to kick off the environment refresh action. The environment refresh will force Delphix to re-examine the Oracle Cluster. This will result in the environment and the dSource being updated based on the cluster's current node, instance, and listener configuration.

DelphixConfiguration Steps to Assign a VDB to a new Cluster Instance

1. Login to the Delphix Management application.
2. In the Oracle RDBMS, add the required number of nodes using the process detailed in the Oracle [RAC documentation](#).
3. In the Delphix Management Engine, select the VDB you want to edit.
4. From the right-side pane select the **Configuration** tab.
5. Refresh the VDB to view new nodes.
6. In the upper right-hand corner, from the **Actions** menu (...) select **Disable**.

The screenshot shows the configuration page for VDBO_8YV. The 'Configuration' tab is selected, and the 'Source' sub-tab is active. The page displays various configuration details for the database, including its unique name, name, version, size, and environment settings. An actions menu is open in the top right corner, showing options like 'Undo Refresh or Rewind', 'Disable', 'Upgrade', 'Migrate', and 'Delete'. The 'Disable' option is highlighted.

Node Na...	Instance ...	Instance ...	Actions
mjr1	1	VDBOMSRD9...	▶

7. In the **Disable** dialog select **Disable**.

The screenshot shows the 'Disable' dialog box. The dialog contains a warning message about disabling the VDB and two checkboxes: 'Force Disable' and 'Provide privileged credentials'. The 'Disable' button is highlighted.


Disabling VDB "VDBO_8YV" will shutdown any running instances, unmount and unexport the filesystem, and prevent any applied policies from executing. In addition, the Engine will no longer actively monitor the database. Are you sure you want to disable this VDB "VDBO_8YV"?

Force Disable


Provide privileged credentials

Cancel Disable

- Click the **pencil** icon next to **Cluster Instances** to be able to add new nodes.



- In the **Include in Cluster** column, check or uncheck the check-box to the left to add or remove a node for your RAC VDB.



- Complete editing and then confirm using the **tick** icon.
- From the Action menu (...), select **Enable**.

Related Topics

- [Adding an Oracle Single Instance or RAC Environment](#)

Using the Double Sync option for Oracle SnapSync

It is possible to request during Linking that a `doubleSync` is performed if `linkNow` is set to true. *(This is available in the GUI for linking starting in 5.0, and for manual Snapshots in 5.3.4.0, see notes below a CLI example).*

When the `doubleSync` option is selected, two SnapSyncs in a row are run, the first one will not even bother to get the archived logs to make the first Snapshot consistent. This means that LogSync will not run during the first SnapSync, or if LogSync is not enabled, the archived logs after the backup of the data will not be retrieved. Once the first SnapSync finishes a second SnapSync is started, starting LogSync if enabled, or taking an archived log backup after the data files have been backed up.

The `doubleSync` option is also available for manually initiated SnapSync, just like `forceFullBackup` option. As a `forceFullBackup` is similar to an initial load, `doubleSync` makes a great companion to `forceFullBackup`, although it is not needed to use both at the same time.

Using the Double SnapSync option via the GUI

1. Login to the **Delphix Management** application.
2. Navigate to the Environment with a Data Source you want to link. Or, from the Datasets page, click the plus icon and select **Add dSource**.
3. In the Add dSource wizard, select the source database with the correct environment user-specified.
4. Enter your login credentials for the source database and click **Verify Credentials**.
If you are linking a mounted standby, see the topics under [Linking Oracle Physical Standby Databases](#) for more information about how the Delphix Engine uses non-SYS login credentials. Click **Next**.
5. Enter a name for your dSource.
6. Select a **Database Group** for the dSource. Adding a dSource to a database group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.

7. In the **Data Management** page select **Show Advanced** and then select **Enable Double SnapSync**. For more information, visit [Data Management Settings for Oracle Data Sources](#).

The screenshot shows the 'Add dSource' configuration page. On the left, a vertical navigation menu includes: Preparation, Source, dSource Configuration, **Data Management** (selected), Policies, Hooks, and Summary. The main content area is titled '1' and contains the following settings:

- Encrypted Linking:**
 - Use system default setting
 - Enable
 - Disable
- Data Load Channels:**
 - Number of Channels:** 2
 - Files per Channel:** 5
- Concurrent files read:** 10
- Block Checking:** Enable
- Level backup:** Enable
- Skip available space check:** Enable
- Double SnapSync:** Enable

8. Assign existing policies to the new dSource. New policies can be created and associated later.
9. Enter any scripts that should be run using the Hooks page.
10. Review the dSource Configuration and Data Management information, and then click Submit.

Linking with doubleSync via the CLI

```
ssh delphix_admin@your engine
delphix > database
delphix database > link
delphix database link*> edit source.operations
delphix database link*> edit postSync
delphix database link*> add
delphix database link*> set command=""
delphix database link*> back; back
delphix database link*> edit preSync
delphix database link*> add
delphix database link*> set command=""
```


```

delphix database link*> back; back; back
delphix database link*> set source.config=XXXX
delphix database link*> set container.name=XXXX
delphix database link*> set container.group=XXXX
delphix database link*> set container.sourcingPolicy.logsyncEnabled=true
delphix database link*> set container.sourcingPolicy.logsyncMode=ARCHIVE_REDO_MODE
delphix database link*> set linkNow=true
delphix database link*> set doubleSync=true
delphix database link*> set dbUser=XXXX
delphix database link*> set environmentUser=XXXX
delphix database link*> set dbCredentials.password=XXXX
delphix database link*>commit


```

Syncing with doubleSync via the GUI

1. Login to the **Delphix Management** application.
2. Navigate to the Datasets page, select the Dataset you want to SnapSync.
3. Click on the arrow next to the camera button.
4. Select **Snapshot with Params.**



5. Select **Double Sync** (and if so desired **Force Full Backup**) from the available parameters.



6. Select **Perform Snapshot.**

Syncing with doubleSync via CLI

```
ssh delphix_admin@yourengine
delphix > database
delphix database > select XXXX
delphix database "XXXX"> sync
delphix database "XXXX"sync *> set doubleSync=true
delphix database "XXXX"sync *>commit
```

Additional Information

 When two Double Sync snapshots are created, then the older snapshot will not be provisionable. Therefore, you should not attempt time flow repair operation as the logs required to make it provisionable, were not retrieved by design. Use the new snapshot for provisioning or refreshing operation.

DoubleSync is not available for multitenant databases.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)
- [Managing Data Sources and Syncing Data](#)

Linking dSources from an Encrypted Oracle Database

This topic describes the behavior of the Delphix Engine when linking to a dSource based on an encrypted Oracle database.

 This topic does not apply to vPDBs.

Beginning with version 10gR2, Oracle supports the encryption of permanent tablespaces using Transparent Data Encryption (TDE). You can link dSources from databases using TDE by following the basic procedure described in [Linking an Oracle Data Source](#). However, in order to provision a VDB from a dSource that is linked to an encrypted database, you must copy wallet files from the physical database in the source environment to the target environment. See [Provisioning a VDB from an Encrypted Oracle Database](#) for more information.

Related Topics

- [Linking an Oracle Data Source](#)
- [Provisioning a VDB from an Encrypted Oracle Database](#)

Detaching and Re-Attaching Oracle dSources

Overview

Each dSource contains metadata that associates it with the source database, as well as the data it has ingested from the source database in the form of snapshots up to that point. It is possible to detach, or unlink, a dSource from its source database. This breaks the association with the source database without affecting the data within the Delphix Engine. Detached dSources and their source databases have these properties:

- A detached dSources can still be used to provision a virtual database (VDB).
- You can re-link the source database as a different dSource.

Prerequisites

Before detaching an Oracle dSource, you must capture the following RMAN configuration:

- Level or SCN based backups
- Data load channel settings: number of channels and files per channel

A dSource can only be attached to a new data source once it has been unlinked. The above RMAN configuration will be removed once the dSource is unlinked.

When attaching an Oracle dSource to a new data source, the new data source must be the same logical database satisfying the following constraints:

- Same dbid
- Same dbname
- Same creation time
- Same resetlogs SCN
- Same resetlogs time
- Same redo stream, where a log must exist with
 - Same sequence
 - Same thread
 - Same end SCN

For Oracle dSources, this procedure can be used to initially link from a standby server that is faster or less disruptive, unlink the dSource, and then attach it to the production server for subsequent incremental SnapSync operations. When you perform the attach operation, you will need the source config name of an unlinked database.

Detaching a dSource

1. Login to the **Delphix Management** application as a user with **OWNER** privileges on the dSource, group, or domain.
2. Click **Manage**.
3. Select **Databases**.
4. Select the **database** you want to unlink or delete.
5. From the **Actions** menu (...) select **Unlink**.
A warning message will appear.
6. Click **Unlink** to confirm.

To detach a dSource via CLI, see [Detaching and Attaching an Oracle dSource via CLI](#).

Rebuilding Source Databases and Using VDBs

In situations where you want to rebuild a source database but retain the existing dSource, you will need to detach the original dSource and create a new one from the rebuilt data source.

1. Detach the dSource as described in the procedure on this page.
2. You cannot attach a dSource with the same name as a dSource that is already attached. If you intend to give the new dSource the same name as the original one, rename the detached dSource.
 - a. At the top of the **Configuration** tab, next to the dSource's name, click the **Edit** (pencil) icon.
 - b. After renaming the dSource, click the green **checkmark**.
3. Create the new dSource from the rebuilt database.

You will now be able to provision VDBs from both the detached dSource and the newly created one, but the detached dSource will only represent the state of the source database prior to being detached.

Attaching a Previously Detached dSource

You can only re-attach databases that represent the same physical database.

Via GUI

To attach a dSource back into the Delphix Engine via GUI, from the **Delphix Management** application:

1. Select **Manage > Datasets**.
2. Select your dSource.
3. From the **Actions** menu (...) select **Link dSource**.

- In the Link dSource dialog, fill in the information corresponding to the new dSource as well as its new environment.

- Click **Link**.

Via CLI

To attach a dSource back into the Delphix Engine via CLI:

- Login to the **Delphix CLI** as `delphix_admin` or a user with OWNER privileges on the dSource, group, or domain.
- Select the dSource by name using `database select <dSource Name>`.
- Run the `attachSource` command.
- Set the source config you want to attach to, using `set attachData.config=<Source Database Unique Name>`. Source configs are named by their database unique name.
- Set any other source configuration operations as you would for a normal link operation.
- Run the `commit` command.

Backup Mode for Attaching Oracle dSources

For Oracle dSources, the SnapSync backup option should be set to **SCN Backup** mode. **Level Backup** mode is based on information stored in the database control file. If the control file of the newly attached database does not contain information about the previous backups, an initial backup will be created. In addition, Block Change Tracking will not be in sync, and the next SnapSync will need to read the entire database to determine which blocks have changed. See [Advanced Data Management Settings for Oracle dSources](#) for more information about Backup Mode.

Related Topics

- [Detaching and Attaching an Oracle dSource via CLI](#)
- [Data Management Settings for Oracle Data Sources](#)

Provisioning from a Replicated Oracle dSource

This topic describes how to provision from a replicated dSource or virtual database (VDB).

The process for provisioning from replicated objects is the same as the typical VDB provisioning process, except that first you need to select the replica namespace containing the replicated object.

Prerequisites

- You must have a dSource or VDB that has been replicated from one Delphix Engine to another, as described in [Replication Overview](#).
- The Delphix Engine containing the replicated dSource or VDB must have a compatible target environment that it can use to provision a VDB from the replicated dSource or VDB.

Procedure

1. On the Delphix Engine containing the replicated dSource or VDB, login to the **Delphix Management** application.
2. In the top menu bar, click **Manage**.
3. Select **Datasets**.
4. From the list of replica namespaces, select the **replica namespace** that contains the dSource or VDB from which you want to provision.
5. The provisioning process is now identical to the process for provisioning standard objects.

Post-Requisites

Once the provisioning job has started, the user interface will automatically display the new VDB in the live system.

Related Topics

- [Replication Overview](#)

Detaching and Re-attaching a PDB dSource from a Data Guard Site

If you want to move the PDB dSource from one Data Guard site to another, then you must detach the PDB from the current environment and attach it to the new environment.

Perform the following procedure to attach or detach a PDB dSource from a Data Guard site.

1. Detach PDB from CDB (a).
 - a. Log into the **Delphix Management** application.
 - b. Select **Manage > Datasets**.
 - c. Select dSource.
 - d. From the **Actions** menu (...), select **Unlink dSource**.
 - e. Click **Unlink** to confirm.
2. If you are going to attach the PDB dSource from a different Data Guard site (Primary or Standby site), perform a fresh refresh of the existing environment to discover that host environment into Delphix. From the **Delphix Management** application:
 - a. Go to **Manage > Environments**.
 - b. If it is already a discovered environment, then from the **Actions(...)** menu, select **Refresh All**. Otherwise, select **Add Environment** to add the environment manually. For more information on adding an environment, see [Adding an Oracle Single Instance or RAC Environment](#). The new database configuration will appear in the Delphix Engine.

 Do not add the discovered PDB as a dSource.

- c. Detach the existing PDB dSource.
 - i. Select the dSource.
 - ii. From the **Actions** menu (...), select **Unlink dSource**.
 - iii. Click **Unlink** to confirm.
- d. Using the Delphix CLI, detach the current CDB datasets associated with the same PDB dSource.

```
dlpx6010 database> select cdb19c1
dlpx6010 database 'cdb19c1'> detachSource
dlpx6010 database 'cdb19c1' detachSource *> ls
Properties
  type: DetachSourceParameters
  source: (required)
dlpx6010 database 'cdb19c1' detachSource *> set source=cdb19c1
dlpx6010 database 'cdb19c1' detachSource *> commit
Dispatched job JOB-16
DB_DETACH_SOURCE job started for "SugoPronto/cdb19c1".
DB_DETACH_SOURCE job for "SugoPronto/cdb19c1" completed successfully.
```

- e. Using the Delphix CLI, attach the new CDB configuration from the new environment/Physical Standby site.

```
dlpx6010 *> database
dlpx6010 database *> select cdb19c1
dlpx6010 database 'cdb19c1'*> attachSource
dlpx6010 database 'cdb19c1' attachSource *> set
attachData.type=OracleAttachData
```

```
dlpx6010 database 'cdb19c1' attachSource *> set
attachData.config=cdb19c1stb
dlpx6010 database 'cdb19c1' attachSource *> set
attachData.environmentUser=OEL7SITDE2/delphix
dlpx6010 database 'cdb19c1' attachSource *> commit;
cdb19c1
Dispatched job JOB-18
DB_ATTACH_SOURCE job started for "SugoPronto/cdb19c1".
Starting validation of attach parameters for database "cdb19c1".
Validation finished for database "cdb19c1".
Obtaining information from source database "SugoPronto/cdb19c1".
The dSource "cdb19c1stb" was successfully linked from source database
"SugoPronto/cdb19c1".
DB_ATTACH_SOURCE job for "SugoPronto/cdb19c1" completed successfully.
```

- f. Using the Delphix Engine user interface, attach the PDB.
 - i. Go to **Manage > Datasets**.
 - ii. Select the PDB from the datasets list.
 - iii. From the **Actions(...)** menu, select **Link dSource**. The Link dSource page appears.
 - iv. Perform the linking of the PDB from the new environment and click **Next**.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)

Working with Oracle Snapshots

This section lists the steps to take a snapshot and delete the same.

Taking a snapshot creates a new snapshot entry in the Oracle dSource's Timeflow. You can use either **Snapshot (Default)** or **Snapshot with Parameters** option for taking the snapshot.

Snapshot (Default)

Perform the following steps to take a snapshot:

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog box, select **Yes**.
6. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the Snapshot you just created.
7. To delete the snapshot, select the snapshot you just created, and from the Actions menu (...), select **Delete Snapshot**.
8. From the **Delete Snapshot** dialog box, select **Delete**.
9. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the snapshot you just deleted.

Snapshot with Parameters

Perform the following steps to take a snapshot:

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the arrow next to the Camera icon and select the **Snapshot with Params...** option.
5. From the Snapshot dialog box, select one of the following:
 - a. **Force Full Backup** - If you select this option, then the Delphix Engine will perform an incremental backup by default. You must select this option only when a full backup is required. Full and Incremental backups consume the same space on the Delphix Engine.
 - b. **Double Sync** - Selecting this option will perform a SnapSync operation as normal. After the first SnapSync is successful, the Engine will immediately perform a second SnapSync without waiting for the Log Files required for the first SnapSync to be made consistent. This is most useful when performing the initial SnapSync (or when "Force Full Backup" is selected) on a very large database that would lead to a large number of archive logs being required to make the SnapSync consistent. Provisioning from a SnapSync that requires excessive recovery is typically time-consuming.
 - c. **Do Not Resume** - If a failure is encountered during the initial SnapSync, the Delphix Engine can resume the SnapSync at a later date. This option will cause the engine to not resume, but rather to start the initial SnapSync over again.
6. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the Snapshot you just created.
7. To delete the snapshot, select the snapshot you just created, and from the Actions menu (...), select **Delete Snapshot**.
8. From the **Delete Snapshot** dialog box, select **Delete**.
9. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the snapshot you just deleted.

Moving the PDB to a New CDB

Perform the following procedure to move the dSource PDB to a new CDB.

1. Follow Oracle steps to unplug your dSource PDB and plug it into your chosen CDB (Refer to the Oracle webpage on [Plugging an Unplugged Pluggable Database](#)). After you have plugged your PDB dSource into the new CDB, discover the host environment that has the new CDB into Delphix, if it is not already discovered. Log into the **Delphix Management** application and do the following:
 - a. Select **Manage > Environments**.
 - b. Select your new CDB environment, then discover the new CDB where your PDB dSource is.
2. If you have already discovered the CDB, then:
 - a. Select **Manage > Environments**.
 - b. Select the dSource Environment where the PDB was plugged.
 - c. **Select the Refresh** icon.
3. When you are ready to attach the PDB dSource, log into the **Delphix Management** application, and do the following:
 - a. Select **Manage > Environments**.
 - b. Select your original PDB dSource Environment.
 - c. **Select the Refresh** icon.
4. To Attach the PDB dSource back into the Delphix Engine, do the following:
 - a. Select **Manage > Datasets**.
 - b. Select your PDB dSource.
 - c. From the **Actions** menu (...), select **Link dSource**.

- d. In the Link dSource dialog, fill in the information corresponding to the new PDB dSource as well as its new environment:

Link dSource ✕

Source Environment

mg-centos-75-oracle-18000-src.dlpxdc.co ▼

Installation

/u01/app/oracle/product/18.0.0.0/dbhome_1 ▼

Database

DBOMSRBBDC6C ▼

Environment User

oracle ▼

Database Username

Database Password

- e. Click **Link**.

Related Topics

- [Linking Data Sources and Syncing Data with Oracle](#)

Provisioning and Managing Virtual Databases with Oracle

Each VDB has its own data management settings, found during the provisioning workflow as well as in the configuration page for that VDB. When you create a VDB, the Delphix Engine copies configuration settings from the dSource and uses them to create the VDB.

This section covers the following topics:

- [Overview of Provisioning Oracle Virtual Databases](#)
- [Configuration Settings for Oracle Virtual Databases](#)
- [Provisioning an Oracle VDB](#)
- [Provisioning an Oracle Virtual Pluggable Database \(vPDB\)](#)
- [Provisioning a TDE-enabled vPDB](#)
- [Provisioning a vPDB from a Non-Multitenant Source](#)
- [Provisioning a VDB from an Encrypted Oracle Database](#)
- [Refreshing and Rewinding a TDE-enabled vPDB](#)
- [Timeflows for RAC Provisioning of VDBs](#)
- [Upgrading an Oracle VDB, Linked CDB, or a vCDB](#)
- [Customizing RAC Instances After Provision](#)
- [Migrating an Oracle VDB](#)
- [Migrating a vPDB](#)
- [Migrating a TDE-enabled vPDB](#)
- [Migrating a vPDB \(from a linked CDB\) to a higher Oracle version \(linked CDB\)](#)
- [V2P with an Oracle VDB](#)
- [Provisioning from a Replicated Oracle VDB](#)
- [Provisioning a TDE-enabled vPDB to a Cluster Target](#)

Overview of Provisioning Oracle Virtual Databases

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment, as described in the overview for [Managing Environments and Hosts](#) and [Requirements for Oracle Hosts and Databases](#).

✔ A Solaris x86 source host is compatible with a Linux x86 target host and vice versa.

From a dSource, you can select a snapshot or point in time to create a VDB. Oracle VDBs each have their own configuration settings as described in [Configuration Settings for Oracle Virtual Databases](#).

Procedure

1. In the Datasets panel on the left-hand side, click the group containing the dSource or VDB from which you want to provision.
2. From the **Timeflow** tab, select a snapshot or point in time to provision from.
3. Click to open the **Provision VDB wizard**, and select a compatible Target Environment for the new Oracle VDB
 - Select Provide Privileged Credentials if you want to use login credentials on the target environment that are different from those associated with the Environment User
 - In the **Advanced** section, you may customize the VDB's configuration settings, file mappings, or custom environment variables.
 - For more information, see [Configuration Settings for Oracle Virtual Databases](#).
 - Select a **Snapshot Policy** for the VDB.
 - Enable **Masked Provisioning** by selecting an option on the Masking page.
 - Enter any operations that should be run in the Hooks page.
 - Review the VDB Configuration and Summary, and then click **Submit**.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the Status tab, or by selecting System and viewing the Jobs page.

Alternatively, you could see this in the Actions Sidebar. When provisioning is complete, the VDB will be included in the group you designated and listed in the Datasets panel.

Provisioning by Snapshot or LogSync

When provisioning by snapshot, you can provision to the start of any particular snapshot, either by time or SCN.

Snapshot Options	Description
Provision by Time	You can provision to the start of any snapshot by selecting that snapshot card from the TimeFlow view or by entering a value in the time entry fields below the snapshot cards. The values you enter will snap to the beginning of the nearest snapshot.

Provision by SCN	To open the SCN entry field select the Actions (...) menu for your selected snapshot and select View by SCN. Here, you can type or paste in the SCN to which you want to provision. After entering a value, it will "snap" to the start of the closest appropriate snapshot.
------------------	--

For more granularity, you can use the LogSync options to provision to any point in time, or to any SCN, within a particular snapshot.

LogSync Options	Description
Provision by Time	Select Open LogSync control to view the time range within that snapshot. Select the clock icon and then select Time to the point to the time from which you want to provision. You can also enter a date and time directly.
Provision by SCN	Select View SCN to view the range of SCNs within that snapshot. You can also type or paste in the specific SCN to which you want to provision. Note that if the SCN does not exist, you will see an error when you provision.

Related Topics

- [Provisioning an Oracle Virtual Pluggable Database \(vPDB\)](#)
- [Configuration Settings for Oracle Virtual Databases](#)
- [Managing Oracle Environments and Hosts](#)
- [Requirements for Oracle Hosts and Databases](#)

Configuration Settings for Oracle Virtual Databases

Each VDB has its own data management settings, found during the provisioning workflow as well as in the configuration page for that VDB. When you create a VDB, the Delphix Engine copies configuration settings from the dSource and uses them to create the VDB.

The following settings are available for Oracle VDBs:

VDB Setting	Explanation
Mount Base	The directory to which Delphix will provide mounted storage. This is where the contents of this VDB will be located.
VDB Configure Parameters	These are customizable parameters of the VDB that you can set either in the provided table or in the text field. This concept is described further below.
Open Database After Provision	<p>Open the VDB immediately after the provision completes. Enabling this option allows you to set the four parameters below:</p> <ul style="list-style-type: none"> • Online Log Size (MB) • Number of Online Log Groups • Enable Archivelog Mode • Generate new DBID for VDB <p>This setting is recommended to tune the Oracle database settings to match your requirements. Each of these four helps Delphix manage the Oracle database optimally.</p> <p>When provisioning an Oracle VDB with the "Open Database After Provision" option unchecked, the subsequent snapshot job fails because the database is not open. The UI will show that the provision failed, but this is not the case. A VDB provisioned this way will exist with the stop/start enable/disable function, but will need to be manually started on every refresh.</p>
Online Log Size (MB)	Also known as the online redo log, this specifies the size of the VDB's redo log. We recommend increasing this value to 1024 MB for improved performance.
Number of Online Log Groups	The number of online log groups to be assigned to this VDB. The number of groups will depend on the configuration of your VDB. For more information, please visit Creating Redo Log Groups and Members .
Enable Archivelog Mode	When you enable this mode, redo logs will be archived instead of overwritten. This creates a backup of all transactions that have occurred in the database so that you can recover to any point in time. These archive logs are useful for recovering a database or updating a standby database.

Generate new DBID for VDB	Create a new Oracle DBID value for this VDB. Note: You can also toggle this option after VDB is created. See Generate a New DBID for Oracle VDBs and Toggle new DBID generation upon refresh options for Oracle VDBs .
Listeners	Select an Oracle listener from the list provided to use with this VDB.
Auto VDB Restart	Enabling this option will automatically restart this VDB whenever its target host is rebooted. Auto VDB Restart can help you recover from host downtime automatically since the Delphix Engine will restart VDBs once they become available.
Configure File Mapping	This option allows you to leverage Oracle’s File Mapping feature with a VDB. Enabling this will add a step to the provisioning wizard to create file mappings.

Using Oracle Multitenant offers different configuration settings, as described in the table below:

Multitenant Setting	Explanation
vCDB Database Parameters	vCDB Database Parameters are Oracle configuration settings you can apply during the provisioning process. You can specify name-value pairs to apply these settings to the vCDB if you are creating one.
Auto vCDB Restart	Enabling this setting will automatically restart this vCDB whenever the target host is rebooted. This will cause all vPDBs in the vCDB to restart as well.

VDB Config Templates

A VDB config template is a list of database configuration parameter names and values that you can save on the Delphix Engine to use at a later time.

Creating a VDB Config Template Via GUI

1. Log into the **Delphix Management** application as an engine administrator.
2. Click **Manage**.
3. Select **VDB Config Templates**.
4. Click the  icon next to the **VDB Config Temp...** and select **New Template** to create a new template.
5. In the **New Template** dialog, enter the name for the new template, parameters that you want to provide, and select the source type from the available options.

6. Click **Create**.

Updating a VDB Config Template Via GUI


1. Log into the **Delphix Management** application as an engine administrator.
2. Click **Manage**.
3. Select **VDB Config Templates**.
4. Select the template from the left-side pane that you need to update.
5. Click on the **pencil** icons to edit an existing VDB template.

6. Click the button to save the changes or click the button to discard the changes that you made.

Apply/Configure a VDB Config Template to a VDB Via GUI

1. Log in to the **Delphix Management** application as an engine administrator.
2. Click **Manage > Datasets**.
3. Select the VDB you want to edit.
4. From the right-side pane select the **Configuration** tab.

5. Click on the **pencil** icon next to 'DATABASE' to be able to edit the database.



The screenshot displays the configuration page for a VDB named 'VDB_CUFMT6SQ'. At the top, there are navigation icons and a title bar. Below the title bar, there are tabs for 'Timeflow', 'Status', and 'Configuration'. Under the 'Configuration' tab, there are sub-tabs for 'Source', 'Policies', 'Masking', and 'Hooks'. The 'DATABASE' section is highlighted with a blue box around a pencil icon, indicating it is the focus of the instruction. The 'ENVIRONMENT' section is also visible below.

DATABASE

Data Path
/work

Data Type
Unstructured Files

Size
12.00KB

Auto VDB Restart @
No

ENVIRONMENT

Name
mg-centos-75-oracle-18000-src.dlpxdc.co

OS
Linux (CentOS)

User
oracle

Timezone
America/New_York,EDT-0400

Additional Mount Points

No items

6. To apply an already existing VDB template, select from the available template in the drop-down under **VDB Config Template**.
Similarly, you can configure the VDB for an already associated template from the available templates.

VDBO_OBA

Timeflow Status Configuration

Source Policies Masking Hooks

Name
VDBO_OBA

SID
VDBOMSRBBD COBA

Version
Oracle 18.3.0.0.0

Size
2.12GB

Archivelog Mode
On

Online Logs
Each log is 1,024 MB, 3 log group(s) per Instance

Auto VDB Restart ⓘ
 Enabled

VDB Config Template

No Template Selected

test

7. Click the button to save the changes or click the button to discard the changes that you made.

You can apply a VDB Config Template to a VDB during the provisioning process, which copies the values from the template. Any subsequent changes to the template will be reflected in the VDB when that VDB is refreshed. During provisioning, you can specify configuration parameters directly or copy them from a VDB Configuration Template. Once set, the Delphix Engine will use these parameters whenever the VDB is refreshed, even if you change the original template. It is important to know, however, that some configuration parameters cannot be customized. In addition, some configuration parameters are stripped out during the provisioning process but are customizable. The list of restricted and customizable parameters can be found below.

Required Parameters

The VDB configuration template must contain at least the following database parameters:

- **compatible** : Required for both Non-Multitenant and Multitenant sources. Must be set with the same value as in the dSource 'compatible' parameter.
- **enable_pluggable_database** : Required for Multitenant sources only. Must be set to TRUE.

Customizable VDB Configure Parameters

The default value for these parameters is cleared during the provisioning process. They are removed from the VDB configuration file unless you set values for them through a VDB Config Template.

- audit_file_dest
- audit_sys_operations
- audit_trail
- background_dump_dest
- core_dump_dest
- db_domain
- diagnostic_dest
- dispatchers
- fast_start_mttr_target
- log_archive_dest_n
- n can be values between 1 - 31
- log_archive_dest_state_n
- n can be values between 1 - 31
- remote_listener
- user_dump_dest

Restricted Parameters

These parameters are restricted for use by the Delphix Engine. Attempting to customize these parameters through the use of a VDB Config Template will cause an error during the provisioning process.

- active_instance_count
- cluster_database
- cluster_database_instances
- cluster_interconnects
- control_files
- db_block_size
- db_create_file_dest
- db_create_online_log_dest_n
- n can be values between 1 - 5
- db_file_name_convert
- db_name
- db_recovery_file_dest
- db_recovery_file_dest_size
- db_unique_name
- dg_broker_config_file1
- dg_broker_config_file2
- dg_broker_start
- fal_client
- fal_server

- instance_name
- instance_number
- local_listener
- log_archive_config
- log_archive_dest
- log_archive_duplex_dest
- log_file_name_convert
- spfile
- standby_archive_dest
- standby_file_management
- tde_configuration
- thread
- undo_tablespace
- wallet_root
- __db_cache_size
- __java_pool_size
- __large_pool_size
- __oracle_base
- __pga_aggregate
- __sga_target
- __shared_io_pool_size
- __shared_pool_size
- __streams_pool_size

Related Topics

- [Automatic VDB Restart on Target Server After Reboot](#)
- [Customizing Oracle VDB Environment Variables](#)
- [Customizing VDB File Mappings](#)

Automatic VDB Restart on Target Server After Reboot

The Delphix Engine now automatically detects whether a target server has been rebooted, and proactively restarts any VDB on that server that was previously up and running. This happens independent of the data platform; it is done as if a target server was restarted and a start command was issued from the Delphix Engine. This feature is compatible with Delphix Self-Service (Jet Stream) ordering dependencies and is limited to non-clustered VDBs.

- i** This feature is not supported for the following data source types:
- Oracle RAC VDBs
 - Oracle 12.1 and older vPDBs that are provisioned into a non-virtual CDB.
 - MSSQL cluster VDBs.

Automatic VDB Restart is supported for following Oracle data sources types:

- Non-Multi-Tenant (non-MT) VDBs
- Oracle 12c and later vPDBs that are provisioned into a virtual CDB. Users can choose to enable or disable automatic restart at the virtual CDB level. There is no individual vPDB automatic restart setting.
- Oracle 12.2 and later vPDBs that are provisioned into a non-virtual CDB. Users can choose to enable or disable individual vPDB's automatic restart.

Please note, the scheduler checks every one minute if any VDB needs to be restarted. The Delphix Engine waits for five minutes after it detects a host has restarted before trying to restart VDBs. The engine will keep trying to restart the VDB for 30 minutes before giving up.

To enable automatic restart, complete the following steps:


1. When provisioning a new VDB in the VDB Provisioning wizard, check the **Auto VDB Restart** box.

The screenshot shows the 'Provision VDB' wizard with the 'Advanced' tab selected. On the left, a vertical progress bar shows steps: Preparation, Source, Provision Point, Target Environment, Target Configuration, **Advanced** (selected), Policies, Masking, Hooks, and Summary. The main content area is titled 'Advanced' and contains the following settings:

- Database After Provision**
 - Open Database After Provision
- Online Log Size (MB)**
 - 1024
- Number of Online Log Groups**
 - 3
- Enable Archivelog Mode
- Generate new DBID for VDB
- Listeners**
 - LISTENER
- Auto VDB Restart** **i**
 - Enabled
- File Mapping**
 - Configure File Mapping
- Custom Environment Variables**

2. Once the VDB has been provisioned, you will be able to turn **Automatic VDB Restart** on.

- a. In the **Datasets** panel, select the **VDB**.
- b. Select the **Configuration** tab.
- c. Select **Source** sub-tab.
- d. Select Database edit.



Related Topics

- [Configuration Settings for Oracle Virtual Databases](#)

Generate a New DBID for Oracle VDBs

Delphix helps to generate a new Oracle internal database identifier (DBID) for VDBs. The DBID is an internal, unique identifier for an Oracle database.

i This feature is not supported for the following data source types:

- Multitenant VDB
- Oracle Live Source

Enable New DBID generation while Provisioning a VDB

To enable generating a new DBID while provisioning a new VDB in the VDB Provisioning wizard, check the **Generate new DBID for VDB** option. This will create a new DBID for the provisioned VDB and subsequent refresh action performed to this VDB will get a different and new DBID.

Provision VDB
✕

- Preparation
- Source
- Provision Point
- Point In Time
- Target Environment
- Target Configuration
- Advanced**
- Policies
- Masking
- Hooks
- Summary

Advanced

Database After Provision

Open Database After Provision

Online Log Size (MB)

Number of Online Log Groups

Enable Archivelog Mode

Generate new DBID for VDB

Listeners

LISTENER

Auto VDB Restart ⓘ

Enabled

File Mapping

Configure File Mapping

Custom Environment Variables

No items

Cancel
Back
Next
Submit

Enable New DBID Generation for an Already Provisioned VDB

If originally the VDB had been provisioned without specifying the **generate new DBID for VDB** option, then you can enable it so that new refreshes will create VDBs with a new DBID (Note the current VDB or previous VDBs will still

have the same DBID as the original database). To accomplish this follow the steps below.

1. In the **Datasets** panel, select the **VDB**.
2. Select the **Configuration** tab.
3. Select **Source** sub-tab.
4. Select Database and click the **pencil** icon to edit.
5. Check the **Generate new DBID upon refresh** option.
6. Refresh the VDB to see the changes.

i If the **Generate new DBID upon refresh** option is enabled, then the Delphix engine will generate a new DBID upon refresh.

The screenshot shows the configuration page for a VDB named **VDBO_MHD**. The interface has a top navigation bar with tabs for **Timeflow**, **Status**, and **Configuration**. Under the **Configuration** tab, there are sub-tabs for **Source**, **Policies**, **Masking**, and **Hosts**. The **Source** sub-tab is active, displaying a **DATABASE** configuration card. This card lists various database properties: **Unique Name** (VDBOMSRE71EE4_MHD), **Name** (VDBO_MHD), **SID** (VDBOMSRE71EMHD), **Version** (Oracle 18.3.0.0.0), **Size** (1.76GB), **Archivelog Mode** (On), and **Online Logs** (Each log is 1,024 MB, 3 log group(s) per Instance). Below these, there are two toggle options: **Auto VDB Restart** (disabled) and **Generate new DBID upon refresh** (disabled). The **Generate new DBID upon refresh** option is highlighted with a blue rectangular box. At the bottom, there is a **VDB Config Template** dropdown menu currently set to **No Template Selected**.

Disable New DBID Generation for an Already Provisioned VDB

You can stop generating the new DBIDs for VDBs that were enabled for this at the time of provisioning or after provisioning. To accomplish this follow the steps below.

1. In the **Datasets** panel, select the **VDB**.
- b. Select the **Configuration** tab.
- c. Select **Source** sub-tab.
4. Select Database and click the **pencil** icon to edit.
5. Uncheck the **Generate new DBID upon refresh** option.
6. This will stop generating a new DBID for subsequent refreshes.
7. Refresh the VDB to reflect the changes.

 If the **Generate new DBID upon refresh** option is disabled, then the Delphix engine will use the same DBID of the parent upon refresh.

Related Topics

[Configuration Settings for Oracle Virtual Databases](#)

Customizing Oracle VDB Environment Variables

Overview

This topic describes how to customize the set of environment variables sourced prior to administering an Oracle virtual database (VDB). Certain Oracle database parameters are sensitive to the environment variables present when you start or administer the database. For this reason, the Delphix Engine allows you to dictate custom environment variables that will be set prior to any administrative action, such as provision, start, stop, rollback, or refresh.

You can specify environment variables by two different means:

- **Name-Value Pair** – A literal variable name and value to be set
- **Environment File** – An environment file to be sourced

Environment variables for Oracle RAC databases might vary in value between cluster nodes. Therefore, environment variable specifications for an Oracle RAC database must specify the cluster node to which they apply.

Setting Custom Environment Variables

Prerequisites

If you are adding any environment variables that are environment files, these files must be accessible on the target environment.

Procedure

1. You can configure custom environment variables in the **Provision Wizard**.
 - a. On the **Target Environment** tab, click **Advanced**.
or
 - b. You can also configure these variables on the **Configuration tab** when the VDB is disabled.
2. Click the **Plus** icon to add an environment variable.
3. Choose a **format** for the environment variable.
 - a. Name-Value Pair
 - i. Enter a **Name** to identify the variable.
 - ii. Enter the variable's **Value**.
 - iii. For Oracle RAC databases, you must also specify the cluster node to which this environment variable applies.
 - b. Environment File
 - i. Enter an absolute path to an environment file on the target environment.
This path can be followed by parameters. Paths and parameters are separated by spaces.
Escaping Spaces
To specify literal spaces, escape them with a backslash ("hello\ world" -> "hello world").
To specify literal backslashes, escape them with a backslash ("foo\\" -> "foo\").
Any other character preceded by a backslash will retain both the backslash and the original character ("\"b" -> "\b").
Escaping is done in order from left to right ("part1\\ part2" -> "part1\" "part2" will be two parameters).
 - ii. For Oracle RAC databases, you must also specify the cluster node to which this environment variable applies.
4. Save the custom environment variables by completing provisioning, or clicking the **Confirm**.
These environment variables will take effect when you start the Oracle VDB.

Environment Variable Denylist

The Delphix Engine denylists the following environment variables; they cannot be set by the user.

- ORACLE_SID
- ORACLE_BASE
- ORACLE_HOME
- CRS_HOME
- ORACLE_UNQNAME
- ORAENV_ASK
- LOGON_STR
- DLPX_SHELL
- SQLPLUS_PLSQL_MODIFIERS
- SQLPLUS_DML_MODIFIERS
- SQLPLUS_DDL_MODIFIERS

If a Name-Value pair has any one of these prohibited environment variables as the name, an error will be raised.

If an environment file sets one of these variables, the Delphix Engine will override this value when the Oracle VDB is started.

The following environment variables will be set before invoking the user-specified script, and thus can be accessed within the script.

- ORACLE_SID
- ORACLE_BASE
- ORACLE_HOME
- CRS_HOME
- ORACLE_UNQNAME

User-Input Sanitation For Environment Variables

For security purposes, user-input provided through the custom environment variables feature retains its literal value when interpreted, including ', ", and \$.

Certain frequently-used environment variables are exempt from this rule and will be interpreted:

- \$ORACLE_HOME
- \$ORACLE_BASE
- \$ORACLE_SID
- \$CRS_HOME
- \$ORACLE_UNQNAME
- \$LD_LIBRARY_PATH
- \$PATH

Caveats

- Environment variables are sourced on provision, start, stop, rollback, and refresh. Custom environment variables are not applicable to V2P.
- Custom environment variables do not propagate to child VDBs and must be set again on provision.
- Custom environment variables do not persist after migration.
On migration of a VDB with custom environment variables, an alert will be raised that the custom

environment variables have been removed from the VDB. In order to view the alert, go to **System > Event Viewer**.

- Custom environment variables are *not* available within scripts executed by VDB Hook Operations.

Related Topics

- [Configuration Settings for Oracle Virtual Databases](#)

Customizing VDB File Mappings

This topic describes how to customize file path mappings when provisioning a virtual database (VDB).

In the VDB provisioning process, it may be necessary to create mappings between files and directories that exist on the source, and files or file directories that exist on the target. An example of this is creating a copy in the target environment of a wallet file for an encrypted database that exists in the source environment.

During VDB provisioning, the archive, datafile, external, script and temp directories are mounted from the Delphix Engine to the <Mount Base>/<Database Unique Name>

You can specify the Mount Base and the Database Unique Name on the Provision VDB card at provision time. But the archive, datafile, external, script and temp mount points are fixed.

- The default Mount Base is set to /mnt/provision
- The default Database Unique Name starts with a capital V and includes the dSource name

Provision VDB

- Preparation
- Source
- Provision Point
- Target Environment
- **Target Configuration**
- Advanced
- Policies
- Masking
- Hooks
- Summary

Target Configuration

Configure the target environment.

Target Group [Add Dataset Group](#)

test

Mount Base

/mnt/provision

VDB CONFIGURATION

VDB Name

VDBO_E0V

Oracle Database Name

VDBO_E0V

Oracle Database Unique Name

VDBOMSRBBD6C_E0V

Oracle SID

VDBOMSRBBDCE0V

VDB Configuration Parameters

Configure VDB Parameters

Mount Base: /mnt/provision

Database Unique Name: Vdbhcp3_35F

Mounts: /mnt/provision/Vdbhcp3_35F/archive

/mnt/provision/Vdbhcp3_35F/datafile

```
/mnt/provision/Vdbhcp3_35F/external
```

```
/mnt/provision/Vdbhcp3_35F/script
```

```
/mnt/provision/Vdbhcp3_35F/temp
```

File Mappings affect everything that follows the mount point directories above. For datafiles, the datafile location from the dSource is used. You can determine the directory structure by querying the Source database:

```
SQL> select name from v$datafile;
```

```
NAME
```

```
-----
```

```
/datafile/dbdhcp3/oradata/dbdhcp3/system01.dbf
```

```
/datafile/dbdhcp3/oradata/dbdhcp3/sysaux01.dbf
```

```
/datafile/dbdhcp3/oradata/dbdhcp3/undotbs01.dbf
```

```
/datafile/dbdhcp3/oradata/dbdhcp3/users01.dbf
```

```
/u03/app/ora11202/product/11.2.0/dbhome_1/dbs/dbv_R2V4.dbf
```

Users have control over the datafiles that were returned by the select name from v\$datafile and also the controlfile location.

Archive log files go directly into /mnt/provision/Vdbhcp3_35F/archive

Tempfiles go directly into /mnt/provision/Vdbhcp3_35F/temp

By default VDB File Mapping would append the above directory structure to the /mnt/provision/Vdbhcp3_35F/datafile directory:

```
/mnt/provision/Vdbhcp3_35F/datafile/datafile/dbdhcp3/oradata/dbdhcp3/
system01.dbf
/mnt/provision/Vdbhcp3_35F/datafile/datafile/dbdhcp3/oradata/dbdhcp3/
sysaux01.dbf
/mnt/provision/Vdbhcp3_35F/datafile/datafile/dbdhcp3/oradata/dbdhcp3/
undotbs01.dbf
/mnt/provision/Vdbhcp3_35F/datafile/datafile/dbdhcp3/oradata/dbdhcp3/
users01.dbf
/mnt/provision/Vdbhcp3_35F/datafile/u03/app/ora11202/product/11.2.0/dbhome_1/
dbs/dbv_R2V4.dbf
```

Pattern Matching Example

You can use pattern matching rules to create full path names for data files and control files.

Pattern matching rules have the form **source-regex-expression-KEY : target-replacement-VALUE**. You can use multiple rules, which are applied successively. Multiple rules with the same source key are allowed.

File Mapping Options

Example 1

For this example, the ultimate goal is to create a VDB that has the following datafile File Mappings:

```

/u03/devvdb3/datafile/sys/system01.dbf
/u03/devvdb3/datafile/sys/sysaux01.dbf
/u03/devvdb3/datafile/ctrl/control01.ctl
/u03/devvdb3/datafile/undo/undotbs01.dbf
/u03/devvdb3/datafile/data/users01.dbf
/u03/devvdb3/datafile/data/dbv_R2V4.dbf

```

You can change the default behavior in a few different ways.

1. Modify the Mount Base.
2. Modify the Database Unique Name.
3. Modify the Mount Base and the Database Unique Name.
4. Modify the name of the individual subdirectories within the Mount Base/Database Unique Name/datafile directory.

Numbers 1-3 can be changed by simply modifying the Mount Base and/or Database Unique Name values when provisioning a VDB:

The datafile locations can be changed in the **Advanced** tab of the **Provision VDB** wizard and scroll down to **File Mapping**:

When you are changing File Mapping, remember that you only need to refer to the part of the path after /Mount Base/Unique Database Name/datafile.

Start by changing all the datafile locations from /u03/devvdb3/datafile/datafile/dbdhcp3/oradata/dbdhcp3 to

/u03/devvdb3/datafile/data by selecting the **+** sign and setting **File Mapping Source File Match** and the **Replacement** and then select **Validate** to see result.

The result will show the new directory structure for all the datafiles and control file relative to the "/Mount Base/Unique Database Name/datafile" directory, but will only show the directories after the datafile directory.

Source File Match	Replacement	Result
/datafile/dbdhcp3/oradata/dbdhcp3	/data	/data/users.dbf

Selecting Validate will show the result of the above mapping, for all the datafiles.

The File Mappings build upon one another, so the first mapping moves almost all the datafiles and the controlfile to /u03//devvdb3/datafile/data.

/u03/app/ora11202/product/11.2.0/dbhome_1/dbs/dbv_R2V4.dbf was not remapped, as it did not contain the directory structure defined in in Source File Match (/datafile/dbdhcp3/oradata/dbdhcp3)

If you wanted the remaining datafiles and controlfile to remain in /data, it would only require the two File Mappings.

The remaining datafiles and controlfile can be relocated by adding to the mapping (select "+" between each Source File Match/Replacement pair).

Source File Match	Replacement
Result	
/u03/app/ora11202/product/11.2.0/dbhome_1/dbs/dbv_R2V4.dbf	/data/
dbv_R2V4.dbf	/data/dbv_R2V4.dbf
/datafile/dbdhcp3/oradata/dbdhcp3/system01.dbf	/sys/
system01.dbf	/sys/system01.dbf
/datafile/dbdhcp3/oradata/dbdhcp3/sysaux01.dbf	/sys/
sysaux01.dbf	/sys/sysaux01.dbf
/datafile/dbdhcp3/oradata/dbdhcp3/undotbs01.dbf	/undo/
undotbs01.dbf	/undo/undotbs01.dbf
/datafile/dbdhcp3/oradata/dbdhcp3/control01.ctl	/ctrl/
control01.ctl	/ctrl/control01.ctl

Select **Validate** between each new entry in order to verify that datafiles are being mapped as expected.

Once all the files are located the way you want them, select **Next** to continue the provision process.

The **Summary** page will show the modifications to **Mount Base** and **Unique Database Name** and will show that **Customized File Mapping** was defined.

After provisioning completes, you can login to the Target server and verify that the datafiles were mapped correctly:

```
SQL> select name from v$datafile;
NAME
-----
--
/u03/devvdb3/datafile/sys/system01.dbf
/u03/devvdb3/datafile/sys/sysaux01.dbf
/u03/devvdb3/datafile/undo/undotbs01.dbf
/u03/devvdb3/datafile/data/users01.dbf
/u03/devvdb3/datafile/data/dbv_R2V4.dbf
```

Example 2

In this example, several rules are applied to the source file path **/app/oracle/oradata/system01.dbf**.

1. Applying the rule **ora:foo** results in: **/app/foocle/foodata/system01.dbf**
2. Applying the rule **foo:bar** results in: **/app/barcle/bardata/system01.dbf**
3. Applying the rule **ora:no** results in an error, because **ora** is no longer found in the pathname.
4. Applying the rule **bar:orane** results in: **/app/oranecler/oraneclerdata/system01.dbf**
5. Applying the rule **ora:yes** results in **/app/yesnewcle/yesnewdata/system01.dbf**

During the pattern matching process, two errors can be generated.


1. **No match for specified mapping rules**
This is the result when no rules match a source file
2. **Invalid regex pattern specified for path mapping**
This is the result of an invalid regex rule mapping

[This topic on the java.regex.util class](#), hosted on [docs.oracle.com](#), shows the regular expression syntax and constructs recognized by the Delphix Engine pattern-matching operations.

Applying VDB File Mappings During the Provisioning Process

1. In the **Target Configuration** tab of the **Provision VDB** wizard, click **Advanced**.
2. Select **Configure File Mapping**.
3. Click **Next**.

- Click the **Plus** icon to add a mapping rule.



- Enter the mapping rule.
- Click **Validate** to see the results of applying the rule.
If not matches are found, you will see an error message.
- Click **Next** to continue with the provisioning process.

Related Topics

- [Configuration Settings for Oracle Virtual Databases](#)

Manually Starting a VDB

When starting an Oracle VDB instance, an initialization file `$ORACLE_BASE_CONFIG/dbs` is used and by default, this directory does not have group write permissions. Previously to Delphix 6.0.10.0, this limitation forced target hosts to either use the instance owner for provisioning or modify the group permissions on this directory. Both options can be seen as a security risk and increase deployment complexity.

With Delphix Engine 6.0.10.0 onwards, a target host no longer requires write permission to the `$ORACLE_BASE_CONFIG/dbs` directory. Delphix Engine will only copy an initialization file `init<ORACLE_SID>.ora`, to this directory if write permissions exist.

All instance startup attempts from Delphix operations will specify an initialization file to use from the Delphix filesystem rather than the default Oracle location. If a VDB needs to be manually started and Delphix was unable to copy the initialization then it must be specified in the startup command. The instance init file is available in the VDB script directory.

You can use the following syntax to manually start the instance:

```
SQL> startup pfile='/mnt/provision/VDBOMSR66E005_610/script/VDBOMSR66E0610/
initVDBOMSR66E0610.ora';
ORACLE instance started.
```

If there is no write permission in the `$ORACLE_BASE_CONFIG/dbs` directory then manually starting the instance will fail:

```
SQL> startup
ORA-01078: failure in processing system parameters
LRM-00109: could not open parameter file '/u01/app/oracle/product/11.2.0.4/dbhome_1/
dbs/initVDBOMSR66E0610.ora'
```

To manually start a VDB, complete the following steps:

1. Ensure that the VDB is enabled and NFS filesystems are mounted. This is required only when the NFS filesystems are offline.
Follow the below steps via CLI.

```
demo source> select VDBOMSR66E005_610
demo source 'VDBOMSR66E005_610'> enable
demo source 'VDBOMSR66E005_610' enable *> set attemptStart=false
demo source 'VDBOMSR66E005_610' enable *> commit
    Dispatched job JOB-47
    SOURCE_ENABLE job started for "VDBOMSR66E005_610".
    Enabling dataset "VDBOMSR66E005_610".
    Exporting storage containers from the Delphix Engine.
    Mounting datasets.
    Dataset "VDBOMSR66E005_610" enabled.
    SOURCE_ENABLE job for "VDBOMSR66E005_610" completed successfully.
```

2. Login to the target host as a Delphix OS user.

3. Locate the VDB scripts directory. This is located at `<toolkit_directory>/Delphix_<engineuid>_<delphix_osuser_id>_<host|cluster>/databases/oracle/<vdb_uniq_name>/<vdb_instance_name>/.`

```
cd /work/Delphix_8e501f827dee_a7a9072fc4fe_2_host/databases/oracle/
VDBOMSR66E005_610/VDBOMSR66E0610
```

4. Run "`setup-oraenv.sh`". This configures the Oracle environment variables such as ORACLE_SID and ORACLE_HOME.
5. Startup the instance by specifying the path to the instance init file in the scripts directory.

```
[oracle@mwrh74-ora11204-tgt VDBOMSR66E0610]$ sqlplus "/as sysdba"

SQL*Plus: Release 11.2.0.4.0 Production on Thu Jul 22 02:40:47 2021

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to an idle instance.

SQL> startup pfile='/work/Delphix_8e501f827dee_a7a9072fc4fe_2_host/databases/
oracle/VDBOMSR66E005_610/VDBOMSR66E0610/initVDBOMSR66E0610.ora';
ORACLE instance started.

Total System Global Area 1068937216 bytes
Fixed Size                2260088 bytes
Variable Size             616563592 bytes
Database Buffers         444596224 bytes
Redo Buffers              5517312 bytes
Database mounted.
Database opened.
SQL>
```

 If Delphix Engine has permission to write to `$ORACLE_BASE_CONFIG/dbs` during provisioning, the instance can be started up with only the "startup" command.

Provisioning an Oracle VDB

This topic describes how to provision a virtual database (VDB) from a dSource or another VDB.

Masked Provisioning

Masked Provisioning is supported on Oracle RAC only when used with "script-based masking".

Prerequisites

- You must have already done one of the following:
 - linked a dSource from a source database, as described in [Linking an Oracle Data Source](#) or
 - created a VDB from which you want to provision another VDB
- You will need to have the correct OS User privileges on the target environment, as described in [Requirements for Oracle Target Hosts and Databases](#).
- If you want to use customized database configuration settings, first create a VDB Config Template as described in [Customizing Oracle VDB Configuration Settings](#).
- If you are creating a VDB from a dSource linked to an encrypted database, make sure you have copied the wallet file to the target environment, as described in [Provisioning a VDB from an Encrypted Oracle Database](#).

VDB Configuration Templates

It is recommended that you always create a VDB Configuration Template prior to provisioning a Virtual Database. This will allow you to customize your parameters for the initial provisioning, and ensure that subsequent changes are reflected on the VDB when refresh and rewind operations are run. See [Customizing Oracle VDB Configuration Settings](#) for more information.

-  Once a VDB has been provisioned, archive logs, flashback logs, redo logs, and datafiles that are not known to be part of the database will be deleted when a SnapSync is performed on the VDB. As a result, make sure that no files are added to the `datafile` file system, as this is used by Delphix to store datafiles. Use the `external` filesystem for other files that are needed to be included in the snapshot that are not datafiles.


For example, creating a second control file on the `datafile` filesystem is prohibited. Such a file will be removed after SnapSync, rendering the VDB unusable. Similarly, block change tracking must not be enabled on a VDB.

Procedure

1. Login to the **Delphix Management** application.
 2. Select **Manage > Datasets**.
 3. In the **Datasets** panel on the left-hand side, click the **group** containing the dSource or VDB from which you want to provision.
 4. Click the **TimeFlow** tab.
 5. Select a **snapshot**.
- For more information on provisioning options.

 You can take a snapshot of the dSource from which to provision. To do so, click the **Camera** icon.

6. **Optional:** Select  to open LogSync timeline.



7. Select  to provision from a point of time within a snapshot. You can select by date or time.
8. Click  and the **Provision VDB** wizard will open:
- For **Oracle Single Instance** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, and **Environment User** will auto-populate with information from the parent.
 - For **Oracle RAC** the fields **Installation Home**, **Database Unique Name**, **SID**, **Database Name**, **Mount Base**, **Instance Number**, **Instance Name** and **Environment User** will auto-populate with information from the parent.

Editable Fields in the VDB Provision Wizard

The following fields are editable:

Installation Home (need to have an additional compatible target)

Database Unique Name

SID

Database Name

Mount Base

Instance Number (RAC Only)

Instance Name (RAC Only)

- If you need to add a new target environment for the VDB, click the green **Plus** icon next to the **Filter Target** field, and follow the instructions in [Adding an Oracle Single Instance or RAC Environment](#).
- Review the information for **Installation Home**, **Database Unique Name**, **SID**, and **Database Name**. Edit as necessary.
- Review the **Mount Base** and **Environment User**. Edit as necessary.
The Environment User must have permission to write to the specified Mount Base, as described

in [Requirements for Oracle Target Hosts and Databases](#). You may also want to create a new writeable directory in the target environment with the correct permissions and use that as the Mount Base for the VDB.

12. Select **Provide Privileged Credentials** if you want to use login credentials on the target environment that are different from those associated with the **Environment User**.
13. Click **Advanced** to customize the VDB online log size and log groups, archivelog mode, local_listener parameter (TCP/IP protocol addresses), additional VDB configuration settings or file mappings, or custom environment variables. If you are provisioning to a target environment that is running a Linux OS, you will need to compare the `SGA_TARGET` configuration parameter with the shared memory size in `/dev/shm`. The shared memory configured on the target host should match the SGA memory target. You can check the Linux OS shared memory size with the command `df -k /dev/shm` and the `SGA_TARGET` configuration parameter by opening the **Advanced** settings and then finding the value for `SGA_TARGET` under **VDB Configuration Templates**.
14. Click **Next**.
15. Select a **Target Group** for the VDB.
16. Enable Auto VDB Restart to allow the VDB to be automatically restarted when the target host reboot is detected by Delphix.

 The **Auto VDB Restart** feature is currently not available for Pluggable Databases, Clusters, or LiveSource.

17. Click **Next**.
18. Select a **Snapshot Policy** for the VDB.
19. Click **Next**.
20. Enter any operations that should be run at Hooks during the provisioning process.
21. Click **Next**.
22. Click **Submit**.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the **Status** tab, or by selecting **Manage/Dashboards** and viewing the **Job History** panel. Alternatively, you could see this in the **Actions Sidebar**. When provisioning is complete, the VDB will be included in the group you designated and listed in the **Datasets** panel. If you select the VDB in the **Datasets** panel and click the **Configuration** tab, you can view information about the database and its Data Management settings.

Provisioning by Snapshot or LogSync

Refer to [Overview of Provisioning Oracle Virtual Databases](#).

Adding or Removing RAC VDB Cluster Node After a VDB is Provisioned

Prerequisites

- Disable the VDB before you edit the instance configuration.

Procedure

After provisioning:

1. Click the **group** containing the VDB.
2. Click the **VDB**.
3. Disable the VDB in order to make changes to instance configuration.
4. Under **Configuration > Source** click the **edit** button and edit the following:
 - a. **Instance Number** for each corresponding instance

b. Instance Name

- c. Check or uncheck the cluster nodes you want for this RAC VDB
5. Click the check button to save changes.
6. Enable the VDB to apply the instance configuration changes.

Related Topics

- [Provisioning a VDB from an Encrypted Oracle Database](#)
- [Requirements for Oracle Hosts and Databases](#)
- [Configuration Settings for Oracle Virtual Databases](#)

Provisioning an Oracle Virtual Pluggable Database (vPDB)

In the Oracle multitenant architecture, there are two main database types: container databases (CDBs) and pluggable databases (PDBs).

 Source CDBs containing Application Container PDBs are currently not supported.

The process of creating virtual pluggable databases is similar to creating non-multitenant virtual databases, with a few additional steps necessary in the multitenant architecture.

 A Solaris x86 source host is compatible with a Linux x86 target host and vice versa.

Delphix supports provisioning Oracle Virtual Pluggable Databases (vPDBs) in two configurations:

1. **Linked Container Databases (Linked CDBs):** Physical CDBs that have been previously provided by the Oracle DBA on the target environment to which Delphix may provision vPDBs. Physical CDBs must be configured and set up specifically for use by Delphix.
2. **Virtual Container Databases (vCDBs):** vCDBs are created by Delphix during the provision workflow for vPDBs. vCDBs currently only support one vPDB per container.

Prerequisites for Provisioning a vPDB to a Linked CDB

There must be a target environment that has an Oracle installation compatible with the Oracle installation of the source CDB and the source PDB. The following database requirements are needed to provision the vPDB to a linked CDB. For more information, see [Requirements for Oracle Hosts and Databases](#).

- Recommend autodiscovery so that the linked CDB can be found. Otherwise, the linked CDB must be manually discovered before provisioning.
- Linked CDB must be running
- Linked CDB must be in ARCHIVELOG mode
- Linked CDB should be using Block Change Tracking (BCT)
- LogSync must be enabled for the Linked CDB.

Provisioning a vPDB to a Linked CDB

For provisioning a vPDB into a Linked CDB, the source CDB and the target CDB must meet the following compatibility requirements:

- The value of the COMPATIBLE parameter in the source CDB must be less than or equal to the value of the COMPATIBLE parameter in the target CDB
- They must have the same endianness
- They must have compatible character sets and national character sets
- They must have the same set of database options installed. For example, if the source CDB is a real application cluster (RAC) database, the target CDB must be a RAC database.

Setting Up Auxillary CDB Parameters

During a vPDB provisioning into a Linked CDB, the Delphix Engine creates a temporary CDB instance on the target environment to recover the vPDB to a consistent state. This temporary CDB will be automatically deleted after the vPDB is provisioned successfully. By default, this temporary CDB is configured to have the same init parameters as the source database. To manage the configuration of the temporary CDB, you must set up [Repository Templates for Oracle Databases](#) during provisioning.

Provisioning a vPDB to a vCDB

When provisioning a vPDB into a vCDB:

- Delphix will provision a virtual CDB (vCDB) from the source CDB to host the vPDB you are about to create.
- There is no need to set up a repository template, as this provision workflow allows a user to directly specify a VDB configuration template, or set individual database parameters.

Procedure

1. In the **Datasets panel**, select an Oracle PDB dSource or a previously provisioned vPDB.
2. From the **Timeflow** tab, select a snapshot or point in time to provision from.
3. Once the Provision wizard is open, you can either provision with a:
 - **Target Linked CDB:** Select an existing container database as the provision target CDB from the **Container Database** drop-down menu of CDBs on that environment.
 - **Target vCDB:** Select the Create a New Container Database checkbox. This will create a new vCDB object in that environment with this new vPDB plugged into it.
- Click **Next** to advance the left-hand pane to the Target Configuration tab, and edit as necessary.
- The Environment User must have permission to write to the specified Mount Base, as described in [Requirements for Oracle Environments and Data](#).
You can also reuse the Delphix toolkit directory, which already exists as the Mount Base, or create a new writable directory in the target environment with the correct permissions and use that as the Mount Base.
- Enter the vPDB Name, Target Group for the vPDB you are about to provision.
- If you selected to create a new target vCDB, configure the vCDB:
 - Enter the vCDB Name, Database Unique Name, and Database Name for the vCDB you are about to provision.
 - Select the Configure vCDB Parameters checkbox if you want to use a VDB Configuration Template. See [Customizing Oracle VDB Configuration Settings](#).
- Click **Next** to advance the left-hand pane to the Advanced tab.
- Select from the available options for your vPDB. The available options are AutovCDB Restart, File Mapping, and custom environment variables.
For more information, see [Customizing VDB File Mappings](#) and [Customizing Oracle VDB Environment Variables](#).
- Click **Next** to advance the left-hand pane to the Policies tab.
- Select the VDB Snapshot policy to be applied to the vPDB.
Select a Retention Policy for the vCDB, if you are provisioning a vCDB.
- Click **Next** to advance the left-hand pane to the Masking tab.
Select the Mask this vPDB checkbox if you want to mask, and select the masking job to be applied.
- Click **Next** to advance the left-hand pane to the Hooks tab, and create any hooks if necessary. For more information, see [Hook Scripts for Automation and Customization](#).
- Review the provisioning summary. Click **Submit** to proceed with provisioning the vPDB.

Related Topics

- [Requirements for Oracle Hosts and Databases](#)
- [Repository Templates for Oracle Databases](#)
- [Configuration Settings for Oracle Virtual Databases](#)
- [Customizing Oracle VDB Environment Variables](#)
- [Customizing VDB File Mappings](#)
- [CLI Cookbook: Provisioning a Virtual PDB from a Non-Multitenant Source Database](#)

Provisioning a TDE-enabled vPDB

Overview


Provisioning a Virtual Pluggable Database (vPDB) to a linked container database first involves using the GUI or CLI to specify the vPDB parameters (such as the vPDB name and target container) along with the snapshot to provision from. Once the provision job is started with these parameters, the Delphix Engine does the following:

1. Mounts the snapshot files on the target host.
2. Creates and opens (in mount mode) the auxiliary container database on the target host, using the snapshot files. The auxiliary container database will have both the CDB and PDB data files from the dSource.
3. Completes recovery to bring the auxiliary container database into a consistent state.
4. Finalizes the state of the auxiliary database and unplugs the vPDB datafiles.
5. Plugs the vPDB into the target database, and opens it in read-write mode for general use.

If the dSource is TDE-enabled, then Delphix will need to perform additional operations to complete the provision of a TDE-enabled vPDB to a TDE-enabled target container (indicated in **red**):

1. Mounts the snapshot files on the target host.
2. **Creates a keystore with the necessary keys to apply encrypted archived log files.**
3. Creates and opens (in mount mode) the auxiliary container database on the target host, using the snapshot files. The auxiliary container database will have both the CDB and PDB data files from the dSource.
4. Completes recovery to bring the auxiliary container database into a consistent state.
5. **Rotates the vPDB master encryption key by generating a new key that is unique to the vPDB and not associated with the source PDB.**
6. **Exports only the newly generated vPDB key to an exported keyfile to enable unplug.**
7. Finalizes the state of the auxiliary database and unplugs the vPDB datafiles.
8. **Imports the vPDB key from the exported keyfile into the target keystore.**
9. Plugs the vPDB into the target database, and opens it in read-write mode for general use.

The following diagram illustrates the provisioning steps.



At each stage of the provisioning process, the keys and exported keyfiles are always on user storage. The exported keyfile is located in the artifact directory, while the auxiliary and target keystores are in the auxiliary keystores directory. Both the artifact directory and auxiliary keystores directory are subdirectories of the TDE keystores root directory, which is either user specified, or if not specified defaults to the toolkit root directory. As for non-TDE-enabled vPDBs, the final vPDB is on Delphix storage while the target linked CDB and its archive logs remain on user storage.

Provisioning a TDE-enabled vPDB

To initiate the provision, Delphix needs the following pieces of information, all of which can be specified in the GUI or CLI:

Parameter	Description	CLI Parameter	Required?
Parent keystore path	Path to a keystore that contains the keys used to encrypt the dSource datafiles. This keystore must be located on the target system. It does not have to be in the location specified by <code>sqlnet.ora</code> or <code>wallet_root</code> , but can be for consistency. It cannot be located on an ASM filesystem. This parameter can be updated if the path is changed.	<code>source.parentTdeKeystorePath</code>	yes
Parent keystore password	Password for the parent keystore. This parameter can be updated if the password is changed.	<code>source.parentTdeKeystorePassword</code>	yes
Exported keyfile secret	When exporting keys to a keyfile from a keystore, Oracle requires a password to be set. Once specified, this cannot be changed for the life of the vPDB.	<code>source.tdeExportedKeyFileSecret</code>	yes
Target keystore password	Password for the target keystore. This parameter can be updated if the password is changed.	<code>sourceconfig.tdeKeystorePassword</code>	yes
Keystores root directory path	Path to a directory on the target host under which all Delphix related TDE artifacts will be created. This includes keystores used by the auxiliary CDB during provisioning and the artifact directories for TDE-enabled vPDBs.	<code>host.oracleParameters.tdeKeystoresRootPath</code>	yes for cluster targets, optional for single instance targets

Adding or Editing the Target Keystore Password

The target keystore password is specified in the GUI in the Databases tab under Environments:

1. Login to the **Delphix Management** application.
2. Click **Manage > Environments**.
3. Click the **Databases** tab for your Environment.

- Next to **TDE Keystore Password** click on the pencil icon to set or update the password.

The screenshot shows the 'Details' tab for an Oracle installation. The 'Databases' section is expanded, showing a list of databases. The 'TDE Keystore Password' field is highlighted with a blue box, and its value is 'Unset'.

INSTALLATION DETAILS	DATABASES
Allow Provisioning Yes	CDOMLOS RB5A3:UNKNOWN:zWmLfU Discover CDB Add dSource
Use as Staging No	Allow Linking Enabled
Oracle Base /u01/app/oracle	Connection Strings + jdbc:oracle:thin:@(DESCRIPTION=(ENABLE=broken)(ADDRESS=(PROTOCOL=tcp)(HOST=10.110.205.202)(PORT=1521)))(CONNECT_DATA=(UR=A)(SID=CDOMLOS RB5A3))) (discovered) ✓
Version 18.3.0.0.0	CDOMSHSRD97A (CDB: CDOMSHSRD97A)
Bits 64	DB Size 2.22GB
Install Group oinstall	TDE Keystore Password Unset

Adding or Editing the Keystores Root Directory Path

The keystores root directory path is specified in the GUI in the Details tab under Environments:

- Log in to the **Delphix Management** application.
- Click **Manage** → **Environments**.
- Click the **Details** tab for your Environment.

4. Next to **Attributes** click on the pencil icon to set or update attributes, including the keystores root.

ATTRIBUTES

Host Address**NFS Addresses****SSH Port****DSP KeyStore Path ⓘ****DSP KeyStore Password ⓘ****DSP KeyStore Alias ⓘ****DSP TrustStore Path ⓘ****DSP TrustStore Password ⓘ****OS**

Linux

Version

CentOS Linux release 8.3.2011

Release

4.18.0-240.el8.x86_64

Time Zone

America/New_York,EST-0500

Total RAM

7.59GB

Processor Type

x86_64

Toolkit Path**TDE Keystores Root**

Provisioning TDE

The remaining three pieces of information are specified during the vPDB provision itself by clicking on the **Enabled** checkbox.

Provision vPDB

- Target Environment
- Target Configuration**
- Advanced
- Policies
- Masking
- Hooks
- Summary

Target Configuration

Configure the target environment.

Target Group Add Dataset Group

test

Mount Base

/mnt/provision

VPDB CONFIGURATION

vPDB Name ⓘ

VCDO_QBH

Oracle Pluggable Database Name ⓘ

VCDO_QBH

Transparent Data Encryption (TDE)

Enabled

Parent Database TDE Keystore Location

Parent Database TDE Keystore Password

TDE Secret for exported keys

Once a TDE-enabled vPDB is provisioned, it can be used the same as a non-TDE-enabled vPDB within Delphix, with the exception of [migrate](#). There are however a few caveats:

- A refresh operation will use the parent keystore for the recovery. If the dSource is rekeyed then the user will need to update the parent keystore with the new keys. Similarly, if the location or password to the parent keystore has changed then they should be updated before the refresh.
- A rewind operation will use the target keystore for the recovery. If the vPDB is rekeyed after it is provisioned, then the rekey will update the target keystore, so it does not need to be updated in Delphix.
- Each disable operation will result in the keys being exported to an exported keyfile in the artifact directory, to be used for a subsequent enable. Refresh and rewind operations will first disable the existing vPDB, so those will also result in a new exported keyfile in the artifact directory.
- Provisioning a second-generation vPDB (vvPDB) from a TDE-enabled vPDB is done in the same manner as a first-generation vPDB, by specifying the TDE parameters during provision. The current keystore for the vPDB can be specified as the parent keystore.

Example Illustrating Movement of Keys

Consider the example illustrated by the vPDB provision above. The vPDB tde_vpdb is provisioned from the dSource CDOMSHSR52CAPDB2 on the VM tde-source18, which is an Oracle database running version 18.11.0. Connecting to this database, we can query v\$encryption_keys to determine the current keys in use by each PDB:

```
SQL> show pdbs
CON_ID CON_NAME          OPEN MODE  RESTRICTED
-----
 2 PDB$SEED              READ ONLY NO
 3 CDOMSHSR52CAPDB1     READ WRITE NO
 4 CDOMSHSR52CAPDB2     READ WRITE NO
 5 CDOMSHSR52CAPDB3     READ WRITE NO
SQL> select con_id, key_id from v$encryption_keys order by con_id;

CON_ID KEY_ID
-----
 1 Ac9MY5kQwU8GvwLYMXImXmMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
 3 AedrXL3aUk9zv+9t7J8ZsVYAAAAAAAAAAAAAAAAAAAAAAAAAAAA
 4 AdDdKibLKU9mv6PDAIvVvH0AAAAAAAAAAAAAAAAAAAAAAAAAAAA
 5 AWdc3ZRaP09Pvw4+2FmLwHAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

The v\$encryption_keys output for this environment shows that there are 3 PDBs within this CDB, all of which are TDE-enabled. In particular, the PDB used for the dSource has a con_id of 4, and an encryption key starting with **AdDdKibL**.

The vPDB tde_vpdb is provisioned to the CDB CDOMSHTG93CF on the VM tde-target18. Connecting to this database, we can again query v\$encryption_keys to determine the keys in use by each PDB:

```
SQL> show pdbs
CON_ID CON_NAME          OPEN MODE  RESTRICTED
-----
 2 PDB$SEED              READ ONLY NO
 3 CDOMSHTG93CFPDB1     READ WRITE NO
 4 CDOMSHTG93CFPDB2     READ WRITE NO
 5 CDOMSHTG93CFPDB3     READ WRITE NO
 6 TDE_VPDB              READ WRITE NO
SQL> select con_id, key_id from v$encryption_keys order by con_id;

CON_ID KEY_ID
-----
 1 AZTc9eKqLk98v8GkQ8/AmaAAAAAAAAAAAAAAAAAAAAAAAAAAAA
 6 AdDdKibLKU9mv6PDAIvVvH0AAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

The key which was originally present in the wallet on the dSource - **AdDdKibL** - has been imported into the target keystore, and has a con_id of 6, which corresponds to the con_id of the vPDB. There are several things to note about the behavior of Oracle and the v\$encryption_keys table:

1. Keys are never deleted from existing keystores by Oracle, only new keys are added. Therefore, if we were to disable the vPDB, which will unmount and unplug it from the CDB, v\$encryption_keys will still show the key as present, with its original con_id, even though it has been unplugged:

```

SQL> show pdbs
CON_ID CON_NAME          OPEN MODE  RESTRICTED
-----
 2 PDB$SEED              READ ONLY NO
 3 CDOMSHTG93CFPDB1    READ WRITE NO
 4 CDOMSHTG93CFPDB2    READ WRITE NO
 5 CDOMSHTG93CFPDB3    READ WRITE NO
SQL> select con_id, key_id from v$encryption_keys order by con_id;
CON_ID KEY_ID
-----
 1 AZTc9eKq1k98v8GkQ8/AmaAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
 6 AdDdKibLKU9mv6PDAIvVvH0AAAAAAAAAAAAAAAAAAAAAAAAAAAAA

```

2. If the wallet is closed for a particular PDB, v\$encryption_keys will not show any entries for that PDB. The wallet status can be determined by querying v\$encryption_wallet.
3. Querying v\$encryption_wallet while the session is attached to CDB\$ROOT will return information about the entire CDB, otherwise, only the keys for the current PDB are returned.

TDE Keystores Root and Artifact Directory

The artifact directory stores the exported keyfiles used during the workflows for TDE-enabled vPDBs. It is located under the keystores root, in the directory `oracle_tde_keystores`. Each TDE-enabled vPDB will have its own directory within the `oracle_tde_keystores` directory, identified by the vPDB name, group name, and a unique identifier, separated by an underscore. If the keystores root directory is not specified, then it defaults to the toolkit directory path.

For example, if the keystores root directory is `/work` (or keystores root is not specified, and the toolkit directory is `/work`), the artifact directory for the vPDB `tde_vpdb` in the group `Encrypted` could be

```
/work/oracle_tde_keystores/tde_vpdb_Encrypted_ce7a47e6-8860-4398-bab0-cf0233fc5e3c
```

Within the artifact directory, there is a subdirectory `exported_keys` which contains within it the exported keyfiles for each timeflow associated with that vPDB. Each time an export is performed, a new exported keyfile is generated with a timestamp. The contents of the artifact directory may change for future releases, but the path to the artifact directory and the naming convention is not anticipated to change.

As the default keystores root directory root is at the same level as the toolkit directory, it will not be overwritten if a host is refreshed through the Delphix Engine and the toolkit updated. It is the customer's responsibility to maintain the artifact directory and ensure that the contents are not lost, as a disk failure could prevent a TDE-enabled vPDB from being accessed. Thus it is recommended that the keystores root directory be on a disk which is regularly backed up.

If a vPDB is moved to a different host (either through the migrate workflow or an enable after a failover, then the artifact directory will need to be copied to the new target host. See [Migrating a TDE-enabled vPDB](#) for details on the manual steps needed for migration.

The artifact directory is not removed when a TDE-enabled vPDB is deleted; the customer can remove it after confirming that the vPDB has been removed (including from any replicated Delphix Engines).

Related Topics

- [Transparent Data Encryption and Delphix](#)
- [Provisioning a TDE-enabled vPDB to a Cluster Target](#)

- [Migrating a TDE-enabled vPDB](#)
- [Provisioning a vPDB from a Non-Multitenant Source](#)
- [Refreshing and Rewinding a TDE-enabled vPDB](#)
- [Locating and Updating the Value of tdeEncryptionKey](#)

Provisioning a vPDB from a Non-Multitenant Source

Delphix supports provisioning a vPDB from a non-multitenant source database. This feature is only available through the API or command-line interface.

This topic describes how to provision a virtual pluggable database (vPDB) from a non-multitenant source database using the command-line interface.

- ⚠** This feature has the following restrictions:
1. Transparent Data Encryption (TDE) is not supported.
 2. The provision point must correspond to a snapshot. Provisioning from a point in time between snapshots is not supported.
 3. The target CDB must be a physical CDB. Virtual CDB targets are not supported.

Prerequisites

Provisioning a vPDB from a non-multitenant source has the following environment requirements:

- Source host with a non-multitenant Oracle 11g or newer source database.
- VDB target host for provisioning a virtual non-multitenant VDB from the source database.
- CDB target host with a running Oracle target version CDB. The target CDB will be automatically linked if it is not already linked.

The target CDB can be a newer Oracle version than the source database (for example, the source is 12.2 and target is 19c). When an upgrade is also required, there are two options for upgrading the database:

- Upgrade Option 1: After provisioning the VDB. This option requires the ability to upgrade to the Oracle target version on the VDB target host.
- Upgrade Option 2: After plugging into the Linked CDB target database.

There are three scripts used during this procedure:

1. **Pre-snapshot Hook on VDB**: This hook will open the database in "read only" mode and issue a call to the `dbms_pdb.describe` procedure to generate an XML file describing the VDB.
2. **Post-snapshot Hook on VDB**: This hook will return the VDB to "read write" mode.
3. **Non-CDB to PDB Script**: This script will run as a hook present on the Linked CDB target host. This should call into the Oracle script `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` to convert the VDB into a PDB. This script should also upgrade the vPDB if doing Upgrade Option 2.

Workflow

1. Link the non-multitenant Oracle 11g or newer source database as a dSource within Delphix.
2. Provision a non-multitenant Oracle VDB from the dSource onto the VDB target host. This will be referred to as the **Golden VDB**.
3. (If using Upgrade Option 1) Upgrade the **Golden VDB** to the Oracle target version: manually upgrade the database and point it to the new Oracle home. This step is only necessary if the source and target Oracle versions are not the same and the data files will not be upgraded when they are converted below.
4. Create a **Pre-snapshot hook** on the **Golden VDB** to open the database in read only mode and issue the `dbms_pdb.describe` procedure call to create an XML file called `delphix_plugin.xml`. The XML file will be used to plug the **Golden VDB** data files into the target CDB. The **Golden VDB** must be open read only during the subsequent snapshot so that the VDB data files do not require recovery when plugging them into the target CDB.

5. (Optional) Create a [Post-snapshot hook](#) on the **Golden VDB** to remove the database from read only mode. The **Golden VDB** can also remain read only.
6. Take a snapshot of the **Golden VDB**.
7. Create a [PDB conversion script](#) named `dx-post-plugin-hook.sh` in the root of the Delphix toolkit directory of the Linked CDB target host. The name of the vPDB being provisioned/converted will be supplied by Delphix as the first parameter to the script when it invokes the script. The VDB data files will already be plugged into the Linked CDB target at the time the script is invoked. The script should do the following:
 - a. (If using Upgrade Option 2) [Upgrade](#) the vPDB data files prior to the conversion.
 - b. Call into `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` and perform any customizations for the multitenant conversion.
8. Select a **snapshot** (point-in-time not supported) on the **Golden VDB** that has the `delphix_plugin.xml` file and provision a virtual PDB to the Linked CDB target. Virtual CDB targets are not supported.
Note: This step can be executed via the **API / CLI only**, and will not be allowed via the Delphix UI.

CLI Procedure to Provision a vPDB from a VDB

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
$ ssh admin@YOUR_ENGINE
```

2. Move to the database provisioning command line object.

```
delphix> database provision
```

3. Set the parameter type to `OracleMultitenantProvisionParameters`.

```
set type=OracleMultitenantProvisionParameters
```

4. Set the login details for the provision and Delphix OS user who is to perform the provision.

```
delphix database provision *> set username=delphix
delphix database provision *> set credential.type>PasswordCredential
delphix database provision *> set credential.password=delphix
```

5. Give the dataset a name.

```
delphix database provision *> set container.name=vpdb
```

6. Place the new dataset in a Group that appears in the Delphix GUI, in this case, the Targets group.

```
delphix database provision *> set container.group=Targets
```

7. Set the destination mount point which Delphix NFS mounts are to be linked to under the virtual PDB. This folder must exist at a file system level on the Linked CDB target host. Do not use single quotes around the mount path.

```
delphix database provision *> set source.mountBase="/mnt/provision"
```

- If automatically restarting the vPDB is not required after a reboot of the Linked CDB target host, set this to option to false. False is possibly a better option given the container database would need to be running prior to any attempt to pull up a vPDB.

```
delphix database provision *> set source.allowAutoVDBRestartOnHostReboot=false
```

- Supply the destination container database name. The container database should already be discovered. This will be where the vPDB will ultimately be placed.

```
delphix database provision *> set sourceConfig.cdbConfig=CDBSTAGE
```

- Name the vPDB. This is what it will appear as in the destination container database.

```
delphix database provision *> set sourceConfig.databaseName=vpdb
```

- Supply the source **Golden VDB** details. In this example, the provision will use the latest snapshot available from the **Golden VDB** as the point in time from which to provision the vPDB. A specific snapshot can also be picked, but an arbitrary point in time is not supported.

```
delphix database provision *> set
timeflowPointParameters.type=TimeflowPointSemantic
delphix database provision *> set timeflowPointParameters.container=gold_vdb
delphix database provision *> set
timeflowPointParameters.location=LATEST_SNAPSHOT
```

- Check that all the settings you require are in place using the "ls" command.

```
delphix database provision *> ls
Properties
  type: OracleMultitenantProvisionParameters
  container:
    type: OracleDatabaseContainer
    name: vpdb (*)
    description: (unset)
    diagnoseNoLoggingFaults: true
    group: Targets (*)
    performanceMode: DISABLED
    preProvisioningEnabled: false
    sourcingPolicy: (unset)
  credential:
    type: PasswordCredential (*)
    password: ***** (*)
  masked: (unset)
  maskingJob: (unset)
  source:
    type: OracleVirtualPdbSource (*)
```

```

name: (unset)
allowAutoVDBRestartOnHostReboot: false (*)
config: (unset)
customEnvVars: (unset)
fileMappingRules: (unset)
LogCollectionEnabled: false
mountBase: /mnt/provision (*)
operations: (unset)
parentTdeKeystorePassword: (unset)
parentTdeKeystorePath: (unset)
tdeExportedKeyFileSecret: (unset)
sourceConfig:
  type: OraclePDBConfig
  cdbConfig: CDBSTAGE (*)
  databaseName: vpdb (*)
  environmentUser: (unset)
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: (unset)
  services: (unset)
timeflowPointParameters:
  type: TimeflowPointSemantic
  container: gold_vdb (*)
  location: LATEST_SNAPSHOT (*)
username: delphix (*)
VirtualCdb: (unset)

```

Operations
defaults

13. Initiate the provision by committing the operation in the CLI.

```

delphix database provision *> commit
vpdb
Dispatched job JOB-333
DB_PROVISION job started for "Targets/vpdb".
Starting provision of virtual PDB database "vpdb" converted from a single tenant
database.
Preparing multitenant container database "CDBSTAGE".
Creating new TimeFlow.
Generating recovery scripts.
Exporting storage.
Preparing XML manifest file prior to plugin.
Plugging in Oracle pluggable database.
Running user-defined post plug hook.
Opening Oracle pluggable database.
Setting OMF destination for Oracle pluggable database.
Creating PDB tempfiles.
Checking Oracle pluggable database plugin violations.
DB_PROVISION job for "Targets/vpdb" completed successfully.

```

- ✔ To refresh the data in the vPDB from production, first, refresh the **Golden VDB** from the dSource, then refresh the vPDB from the new snapshot in the **Golden VDB**.

- ⓘ There are some workflow customizations required for RAC databases:
1. The PDB conversion script must be in the root of the Delphix toolkit directory for all the target CDB RAC instances.
 2. The **Golden VDB** Pre-Snapshot hook, as provided below, will not work in a clustered (RAC) environment with more than one active instance because it only shuts down the local instance. `dbms_pdb.describe` will not execute while an instance is open read-write. The workarounds are:
 - a. Provision the **Golden VDB** as single-instance, either by provisioning to a non-RAC target or by provisioning to a RAC target with only one active instance. The sample hook will work in this case.
 - b. Write a customized pre-snapshot hook that shuts down all instances, restarts only one instance in read-only mode, and runs `dbms_pdb.describe`.
 - c. Manually perform the actions of the hook: shutdown the **Golden VDB**, restart one of the instances in read-only mode and then run `dbms_pdb.describe`.

Sample Scripts

Golden VDB Pre-Snapshot Hook

Restarts the source VDB in read only mode and runs `dbms_pdb.describe` to generate an XML file describing the VDB. The XML file will be used to plug the VDB into the Linked CDB target. The target for the XML file must be `$DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/delphix_plugin.xml`.

```
#!/bin/sh

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/presnapshot.out
replace
shutdown immediate
startup mount
alter database open read only;
exec dbms_pdb.describe(pdb_descr_file=>'$DELPHIX_MOUNT_PATH/
$DELPHIX_DATABASE_UNIQUE_NAME/datafile/delphix_plugin.xml');
exit;
EOF
```

Golden VDB Post-Snapshot Hook

This is only necessary if the VDB should not be left in read-only mode after the snapshot.

```
#!/bin/sh

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/postsnapshot.out
replace
  shutdown immediate
  startup
  exit;
EOF
```

PDB Conversion Script

The script should be named `dx-post-plug-hook.sh` and reside in the root of the Delphix toolkit directory of the Linked CDB target host. Delphix will supply the name of the PDB being provisioned/converted as the first parameter.

The VDB datafiles will have already been plugged into the target CDB at the time the script is invoked and the virtual PDB will be in the mounted (not open) state. The PDB conversion script should return with the virtual PDB in either the mounted or open (not restricted) state. Delphix does not enforce a time-out for the script.

```
#!/bin/sh

DELPHIX_PDB_NAME=$1
SCRIPT_DIR="$( cd "$( dirname "$0" )" && pwd )"
CONVERT_LOGFILE=$SCRIPT_DIR/$DELPHIX_PDB_NAME-pdbconvert.log

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $CONVERT_LOGFILE replace
  alter session set container=$DELPHIX_PDB_NAME;
  @?/rdbms/admin/noncdb_to_pdb.sql
  exit;
EOF
```

PDB Upgrade Script

The following script will upgrade the vPDB. Use a wrapper that runs both this script and the prior conversion script (or combine the two in a single script) if doing both an upgrade and a conversion.

```
#!/bin/sh

DELPHIX_PDB_NAME=$1
SCRIPT_DIR="$( cd "$( dirname "$0" )" && pwd )"
UPGRADE_LOGFILE=$SCRIPT_DIR/$DELPHIX_PDB_NAME-dx-post-plug-upgrade.log
UPGRADE_LOGDIR=$SCRIPT_DIR/$DELPHIX_PDB_NAME-upgrade

mkdir $UPGRADE_LOGDIR
cd $ORACLE_HOME/rdbms/admin
switches="-c '$DELPHIX_PDB_NAME' -l $UPGRADE_LOGDIR"
$ORACLE_HOME/perl/bin/perl catctl.pl $switches catupgrd.sql &>> $UPGRADE_LOGFILE
```

Related Topics

- [Transparent Data Encryption and Delphix](#)
- [Provisioning a TDE-enabled vPDB](#)
- [Refreshing and Rewinding a TDE-enabled vPDB](#)
- [Migrating a TDE-enabled vPDB](#)

Provisioning a VDB from an Encrypted Oracle Database

Overview

This topic describes how to provision a VDB from an encrypted database. The Delphix Engine supports provisioning from a dSource linked to a physical database that has been encrypted with Oracle's Transparent Database Encryption (TDE), which can be used to encrypt columns or tablespaces.

Note

The Delphix engine does not support provisioning from a dSource with an encrypted system tablespace in a non-multitenant configuration.

Provisioning a VDB from an encrypted dSource requires an auto-open wallet setup in the target environment, because the provisioning process requires the master key to be stored in the wallet file. On the dSource, export the keys and copy the export file (both `ewallet.p12` and `cwallet.sso`) to the VDB server. On the VDB server, create an empty wallet and import the keys into the wallet, then set the key to be used.

When provisioning a VDB from an encrypted dSource, if the target environment has other databases that also use TDE, each database should use a different wallet. This also includes a scenario where the VDB has been provisioned back to the same environment as the encrypted dSource. Please check Oracle documentation on how to set up different wallet locations for different databases. For example, use `$ORACLE_SID` in the `DIRECTORY` clause of the `ENCRYPTION_WALLET_LOCATION` parameter in `sqlnet.ora`.

```
ENCRYPTION_WALLET_LOCATION=(SOURCE=(METHOD=FILE) (METHOD_DATA=(DIRECTORY=/opt/oracle/wallets/$ORACLE_SID)))
```

Procedure

1. Check for any encrypted columns or tablespaces on the source database by using these commands:

```
SELECT t.name name, e.encryptionalg algorithm FROM v$tablespace t,
v$encrypted_tablespaces e

WHERE t.ts# = e.ts# and upper(e.encryptedts) = 'YES';
```

2. Copy wallet files from the source database to the target environment, and then configure the `sqlnet.ora` file on the target to point to the directory where the wallet is located.

```
$ more sqlnet.ora

ENCRYPTION_WALLET_LOCATION=(SOURCE(METHOD=file) (METHOD_DATA=(DIRECTORY=/opt/oracle/oradata/nf/wallet)))
```

3. If the source database does not use the auto-open wallet, create the auto-open wallet at the target environment.

```
$ orapki wallet create -wallet /opt/oracle/oradata/nf/wallet -auto_login [-pwd  
password]
```


4. Proceed with provisioning the VDB as described in [Provisioning an Oracle VDB](#).

Related Topics

- [Provisioning an Oracle VDB](#)
- [Requirements for Oracle Hosts and Databases](#)

Refreshing and Rewinding a TDE-enabled vPDB

Just like a non-TDE-enabled vPDB, a TDE-enabled vPDB can be refreshed from the dSource or rewound to a previous snapshot or point in time. In each case, no additional manual steps or input from the user is required. The first step of a refresh or rewind operation is to disable the existing vPDB, which will result in a new keyfile exported to the artifact directory. The appropriate snapshot files are then mounted for the auxiliary database so that it can be recovered and brought to a consistent state. Since the vPDB is TDE-enabled, a keystore is needed for the recover operation. For a refresh, the Delphix Engine will use the parent keystore, and for a rewind, the Delphix Engine will use the target keystore, as shown below.



Key Rotation

There are two potential places for keys to be rotated in a vPDB environment:

1. **dSource:** If the dSource keys are rotated and a new snapshot taken with the new key, the customer is responsible for updating the parent keystore before refreshing from the later snapshot encrypted with the new key. The parent keystore would then contain both the new key and the original keys.
2. **Target:** If the target CDB keys are rotated, the target keystore will be updated. This is why the Delphix Engine uses the target keystore for rewind operations.

In either scenario, the keystore used for recovery will contain the current and all prior keys used to encrypt the datafiles and archive logs, for both the vPDB and CDB used in the auxiliary container.

vPDB Encryption Key Management

During the provisioning process of a TDE-enabled vPDB, Delphix generates a unique encryption key for the vPDB. This unique key is not associated with the parent keystore to ensure that no keys from the parent are imported by the target. During refresh and rewind operations, Delphix reuses that key after recovery has finished. It is possible to customize the key that is used by updating the `tdeKeyIdentifier` parameter of the source via the CLI. If a valid `key_id` is entered for a key that is already present in the keystore, that key will be used as the active encryption

key for the vPDB at the end of refresh/rewind. If the field is unset, Delphix will generate a new encryption key for the vPDB to be used from that point onward. See the CLI steps for [Locating and Updating the Value of tdeEncryptionKey](#).

Related Topics

- [Provisioning a TDE-enabled vPDB](#)
- [Provisioning a vPDB from a Non-Multitenant Source](#)
- [Transparent Data Encryption and Delphix](#)
- [Migrating a TDE-enabled vPDB](#)
- [Locating and Updating the Value of tdeEncryptionKey](#)

Timeflows for RAC Provisioning of VDBs

Overview

This topic describes special considerations when provisioning by timestamp from a RAC time flow.

Timestamps in Oracle RAC time flows can be imprecise because of time skew among the hosts in a RAC configuration. The time stamps will generally track the host with the fastest clock. For this reason, provisioning by a timestamp may not leave the VDB provisioned at the exact time desired. The provision by SCN should be used if more fine-grained control is required when provisioning.

Introduction

The Delphix Engine provides the ability to link to an external database by creating a dSource within the Delphix system. Once linked, the Delphix Engine maintains a complete history of the database as part of a Timeflow, limited by the retention policies configured by the administrator. From any time within that Timeflow, you can provision a virtual database (VDB) from the Delphix Engine. This Timeflow is maintained through the use of SnapSync and LogSync.

The SnapSync operation pulls over the complete data set of the external database during the initial load. Subsequent SnapSync operations pull and store only incremental changes. At the end of each SnapSync operation, the Delphix Engine creates a snapshot that serves as the base point for provisioning operations. In addition, LogSync periodically connects to the host(s) running the source database and pulls over any log files associated with the database. These log files are stored separately from the SnapSync data and are used to provision from points in between SnapSync snapshots. Usually, SnapSync operates against a live database with changes actively being made to it. Hence the data that it pulls over is “fuzzy” and logs must be applied to the data to make it consistent and provisionable. If LogSync is enabled, SnapSync relies on it to copy the logs over. If LogSync is not enabled, SnapSync copies the logs itself. Occasionally, LogSync or SnapSync is not able to retrieve one or more log files from the database. This creates a break in the Timeflow or can prevent a snapshot from being provisioned. To remedy this situation, the Delphix Engine has tools to repair, or patch, a snapshot and the Timeflow.

Snapshot Repair

ASM

The steps below do not apply if your archive logs are stored on ASM. If they are stored on ASM, you must move the archived logs to a supported filesystem directory.

When missing log files prevent the Delphix Engine from provisioning a snapshot, you can use the Delphix Management application to identify the missing logs and repair the snapshot. The Delphix Engine will generate a fault whenever missing logs prevent a snapshot from being provisionable. The fault will likely have the title "Cannot provision database from snapshot" and will contain a description of the cause. The common causes are:

- Logs were deleted/moved/archived from the database before the Delphix Engine could retrieve them. In this case, the archive log retention policy on the source database may be too aggressive. Use the GUI snapshot repair tool to fetch the logs.
- LogSync is still fetching the logs. SnapSync is relying on LogSync to fetch the logs needed to make the snapshot consistent. SnapSync normally will wait up to 15 minutes for LogSync to fetch the logs. If LogSync has not fetched the logs by then, SnapSync will generate a fault and finish. The best course of action, in this case, maybe to wait for LogSync to fetch the logs.
- The source database is a physical standby in real-time apply mode. The changes described in the current online log of the database are needed to make the snapshot consistent. LogSync cannot retrieve the log


until it is archived, and SnapSync cannot force the log to be archived because the source database is a physical standby. Force a log switch on the primary database or wait until the log is naturally archived.

Below is a screenshot of a snapshot with missing logs. Clicking on the snapshot causes the list of missing log(s) to appear. In this example, log sequence 404 is missing.


The screenshot displays the Delphix Management interface. The top navigation bar includes 'Manage', 'Resources', 'System', and 'Help'. The main content area shows a dataset named 'dbdhcp3-ds-bbdhcp.dlpxdc.co'. Under the 'Timeflow' tab, a snapshot from 'JAN 30, 2019' is shown at '5:30 PM'. A blue box highlights the 'Missing Archive Logs' section, listing '404'. To the right, the 'Dataset Details' panel shows 'STORAGE USED' (1 Snapshot, 21.80GB Total Engine Storage) and 'LOCATION & LINEAGE' (Environment: ds-bbdhcp.dlpxdc.co).

Snapshot with missing logs

If the snapshot can be repaired by fetching the logs from the source database, then you can use the GUI snapshot repair tool to fetch the logs. Hovering over the snapshot exposes more options ("...") which can then be clicked to show the **Repair Missing Archive Logs** option. Clicking on the option starts the repair tool.



Selecting the repair tool



Repair tool

To use the snapshot repair tool, as seen above:

1. Enter a **Hostname**. This should be the host from which to retrieve the log(s).

2. Enter a **Username** and **Password**. These should be the credentials for a user who can read the archived log file(s). The user credentials are optional if the host and user credentials have already been added to the Delphix Engine.
3. Enter a **File Path**. This should be the name of the directory containing the missing log(s).

If more than one file is missing, they should all exist in the directory specified by **File Path**. The tool will read every file in the **File Path** directory, so it is best that it only contains the files that are to be retrieved.

Timeflow Patching

When missing log files cause a break in the Timeflow, you can use the command-line interface (CLI) to identify the missing logs and patch the Timeflow. The Delphix Engine will generate a fault whenever there are missing logs on a portion of the Timeflow. The fault will likely have the title “Cannot provision a database from a portion of Timeflow” and will contain a description of the cause. The most common cause is an overly aggressive archive log retention policy on the source database causing a log to be deleted before LogSync can fetch it. Other faults can also be generated by describing the specific errors encountered when fetching the log(s).

You can use the CLI to list the missing logs and patch the Timeflow. The following CLI Cookbook entry demonstrates how to do this: [CLI Cookbook: Repairing a Timeflow](#).

If you delete or move archivelogs from the source database that are not needed for a snapshot, you still may need to repair the TimeFlow to provision using LogSync. In this case, an icon will not be visible on the TimeFlow tabs. This means you cannot repair the TimeFlow in the GUI. However, you can still repair it using the CLI.

Related Topics

- [CLI Cookbook: Repairing a Timeflow](#)
- [Provisioning and Managing Virtual Databases with Oracle](#)

Upgrading an Oracle VDB, Linked CDB, or a vCDB

Prerequisites

Upgrading a VDB involves Oracle level activity and Delphix GUI/CLI level activity. The Oracle activity, to be performed first, involves upgrading/patching current oracle homes, installing new oracle homes, and managing init.ora/spfile.ora files. The Delphix GUI/CLI activity, to be performed second, involves rediscovering new oracle installations, giving a VDB a new oracle home, and updating the oracle grid home. These are discussed in detail below.

 During an Oracle upgrade, refreshing an Environment will generally discover new Oracle installations, allowing for a virtual database (VDB) or dSource upgrade to be handled through the UI. However, if the Oracle Grid home was changed, due to an Oracle upgrade, the crsClusterHome parameter will need to be updated manually through the command line (CLI).

Limitations

Currently, it is not possible to convert an existing VDB into a vPDB.

PSU/Oracle Upgrade Procedure

Normally a PSU or Oracle upgrade will have both binary changes and some scripts to run on the database side as well.

There are 3 ways to apply a PSU/Oracle upgrade:

- A) Apply to existing ORACLE_HOME. You must be on Delphix version 4.1.x or higher to do this.
- B) Create a new ORACLE_HOME (could clone the existing one) and then apply the PSU to the new ORACLE_HOME
- C) After a dSource is upgraded, use refresh on the Timeflow tab to upgrade the VDB

Follow Oracle documentation and run the appropriate script(s) and/or steps on the databases using those ORACLE_HOMEs. In option B, stop the instance using the old ORACLE_HOME, then restart the instance with the new ORACLE_HOME from the command line as normal.

Applying to an existing ORACLE_HOME

1. Following Oracle documentation, patch the ORACLE_HOME, then the database for the VDB(s).
2. Refresh the environment the VDBs are on in the Delphix Management application.

Creating a new ORACLE_HOME

1. Refresh the environment from the Delphix Management application. Verify that the new ORACLE_HOME is picked up and displayed in the **Databases** tab as an ORACLE Installation.
2. Stop the VDB instance (old ORACLE_HOME) using Oracle tools. Do not use the Delphix VDB **stop** operation as the VDB should be stopped outside of Delphix.
3. Export ORACLE_HOME=(newORACLE_HOME). Follow Oracle documentation to patch the database.
4. Copy the init.ora for that VDB in this new \$ORACLE_HOME/dbs directory. The delphix_os user will need the write permissions to this directory.
5. If there are any database parameter changes, update the spfile located on the Delphix mount base with those values.
6. Navigate to the **Datasets** view
7. Expand the group(s) containing all non-multitenant and multitenant VDBs.
8. Click the **Configuration** tab.

9. From the Actions menu (...) select **Upgrade** to switch the ORACLE_INSTALLATION to the new one.
10. If the database is a linked CDB or vCDB, go to any child vPDBs and verify that the **Repository** and/or **Version** has been updated under the **configuration** tab.

Using Refresh

1. Refresh the environment from the Delphix Management application.
2. Verify that the new ORACLE_HOME is picked up and displayed in the **Databases** tab of the **Environments** screen as an ORACLE Installation.
3. On the **VDB Configuration tab**, click the **stop** icon to stop the VDB.
4. From the Actions menu (...) select **upgrade** to switch the ORACLE_INSTALLATION to the new upgrade version same as the dSource.
5. Navigate to the **Datasets** view, select the VDB, and then select the **Timeflow** tab.
6. Click the **Refresh** button.
7. Select a new **snapshot** from the dSource that was taken after the dSource was upgraded. (The database version is on the snapshot card.)

✓ Updating the Oracle User after an Upgrade

There may be cases when you upgrade the Oracle home and the Oracle User (who owns the binary) is a different user than the previous Oracle User. You will then need to update the Oracle User for each environment, and then re-connect each VDB to the upgraded Oracle home using the new Oracle User.

The new Oracle User must be in the same OS group (for example, dba or oinstall) as the previous one.

1. Login to the Delphix Management application using delphix_admin credentials.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **environment** where you want to add the user.
5. Next to **Environment Users**, click the **Pencil** icon to add the new user.
6. Set the new user as the **default** user.
7. Follow the procedure to upgrade VDBs described in this topic.

Linked CDB Upgrade Procedure

For vPDBs in the linked CDBs, there are two ways to upgrade the linked CDB:

- A) Perform the Oracle upgrade of the current target CDB.
- B) Create an entirely new target CDB of the higher version.

Performing the Oracle Upgrade of the Current Target CDB

1. Perform the Oracle upgrade of the current target CDB.
2. Login to the Delphix Management application.
3. Navigate to **Manage > Datasets** and select the target CDB.
4. From the **Actions** menu (...), select **Upgrade** to switch the ORACLE_INSTALLATION to the new one.
5. Click Upgrade.
6. Under the **Configuration** tab, verify that the **Repository** and/or **Version** has been updated.

Creating a New Target CDB

You can first disable the vPDB, then use the migrate vPDB feature to select a new container database.

Procedure

1. Login to the Delphix Management application.
2. Navigate to **Manage > Datasets**.
3. Select the vPDB you want to migrate.
4. From the **Actions** menu (...) select **Disable**.
5. Click **Disable** to confirm.
6. From the **Actions** menu (...) select **Migrate**.
7. Select the new container database for the vPDB, the user for that environment, and the database installation where the container database of the vPDB will reside.
8. Click the **Migrate** option to confirm your selections.
9. Manually copy the exported keys to the target host's toolkit directory.
10. From the **Actions** menu (...) select **Enable**.
11. Click **Enable** to confirm.
Your vPDB will re-start in the new environment, and you can continue to work with it as you would any other vPDB.

Related Topics

- [Provisioning an Oracle VDB](#)
- [Configuration Settings for Oracle Virtual Databases](#)

Customizing RAC Instances After Provision

Provisioning a RAC VDB requires users to select which RAC node to be included for this VDB, as well as the instance number and instance name of each RAC node that runs the RAC VDB. At any time after the provision, there may be a need to downsize or increase the number of instances of the RAC VDB or reassign instance number and instance names to individual VDB RAC node.



Oracle Pluggable Database

The current release does not support the customization of RAC vPDB instances after provision.




Note

If you make any changes to the VDB's configuration on the engine either by changing the VDB's configParams using CLI or VDB configuration template that was used to provision it, the Delphix engine will push out these changes when the VDB is next enabled after the RAC instance changes are made.

Procedure

The following steps illustrate how you can customize your RAC VDB instance configuration after provision:

1. Login to your Delphix Management application.
2. Select the VDB you want to edit.
3. From the right-side pane select the **Configuration** tab.
4. From the Action menu (...), select **Disable**.



5. Click on the **pencil** icon next to 'CLUSTER INSTANCES' to be able to edit your instance names.

- Click on the name of the instance you want to alter and edit it to suit.

The screenshot shows the Delphix Management console interface. The main content area displays the configuration for a VDB instance named 'VDBO_5N6'. The 'CLUSTER INSTANCES' section contains a table with the following data:

Include in Cluster	Node Name	Instance Number	Instance Name
<input checked="" type="checkbox"/>	rac8k826311	1	VDBOMSRB331B3...
<input type="checkbox"/>	rac8k826312	2	v2

A modal dialog is open over this table, allowing for editing the cluster nodes. The dialog contains a table with a checked checkbox and a blue checkmark button.

- Check or uncheck the check-box to the left to indicate if you want to add or remove a cluster node for your RAC VDB.
- Complete editing and then confirm using the tick icon.
- From the Action menu (...), select **Enable**.

i Disabling a VDB will result in the VDB Instances being shut down and the NFS mounts for that VDB presented from Delphix being dismounted.

Related Topics

- [Timeflows for RAC Provisioning of VDBs](#)

Migrating an Oracle VDB

This topic describes how to migrate a Virtual Database (VDB) from one target environment to another.

There may be situations in which you want to migrate a virtual database to a new target environment, for example when upgrading the host on which the VDB resides, or as part of a general data center migration. This is easily accomplished by first disabling the database, then using the Migrate VDB feature to select a new target environment.

Prerequisites

- You should have already set up a new target environment that is compatible with the VDB that you want to migrate.
- You cannot migrate a single instance of Oracle VDB to a RAC environment and vice versa. An additional reconfiguration is needed when converting a single instance to RAC that is only performed during a VDB provision. Instead, you should provision a new VDB.

Procedure

1. Login to your Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **VDB** you want to migrate.
5. From the Actions menu (...) select **Disable**.
6. Click **Disable** to confirm.
When the VDB is disabled, its icon will turn grey.
7. From the Actions menu (...) select **Migrate**.
8. Select the new target environment for the VDB, the user for that environment, and the database installation where the VDB will reside.

Migrate ✕

Current Environment

Database dbdh_6U8	Environment Starfire
Installation /u01/app/ora10205/product/10.2.0/db_1	User ora10205

New Environment

Environment

Starfire
▾

User

ora10205
▾

Installation

/u01/app/ora10205/product/10.2.0/db_1
▾

If the new installation is not in the current list, you can try to discover it by refreshing the environment. If the new installation is not discovered automatically, you can add it manually.

Cancel
Migrate

9. Select **Migrate** to confirm your selections.
10. From the Actions menu (...) select the **Enable**
11. Click **Enable** to confirm.

Within a few minutes, your VDB will re-start in the new environment, and you can continue to work with it as you would any other VDB.

Related Topics

- [Provisioning and Managing Virtual Databases with Oracle](#)

Migrating a vPDB

There may be situations in which you want to migrate a virtual pluggable database (vPDB) to a new container database on the same or a different target environment, for example when upgrading the host on which the vPDB resides, or as part of a general data center migration. This is easily accomplished by first disabling the vPDB, then using the Migrate vPDB feature to select a new container database.

i The current release only supports migrating from a vPDB in a virtual CDB, to another linked or virtual CDB using the CLI. Attempts to perform this migration via GUI will fail with the following error:

Error
The operation could not be completed because the virtual database “vPDB1” is disabled.

Error Code
exception.db.genericvdb.disabled

Pre-requisites

You should already set up and have Delphix discover a container database in the same environment as the vPDB currently is or from an environment to which the vPDB will be migrated to.

Procedure

Login to your Delphix Management application.

1. Click **Manage**.
2. Select **Datasets**.
3. Select the **vPDB** you want to migrate.
4. From the Actions menu (...) select **Disable**.
5. Click **Disable** to confirm.
6. From the Actions menu (...) select **Migrate**.
7. Select the new **container database** for the vPDB, the **user** for that environment, and the **database installation** where the container database of the vPDB will reside.
8. Click the **Migrate** to confirm your selections.
9. Manually copy the exported keys to the target hosts toolkit directory.
When migrating a TDE enabled vPDB, the disable step results in the vPDB keys being exported to the toolkit. Enabling the vPDB in a different target host will result in plugin violations, therefore when migrating to a new host, the user must manually copy the exported keys to the toolkit directory on the target host.
10. From the Actions menu (...) select **Enable**.
11. Click **Enable** to confirm.
Your vPDB will re-start in the new environment, and you can continue to work with it as you would any other vPDB.

Related Topics


- [Provision an Oracle VDB](#)
- [Configuration Settings for Oracle Virtual Databases](#)

Migrating a TDE-enabled vPDB

Migrating a vPDB from one container database to another involves the following steps:

1. Disable the vPDB.
2. Migrate.
3. For a TDE-enabled vPDB, there are additional steps that are required before the enable. The steps are necessary because the keystore and exported keyfiles are not present on Delphix storage. These steps are:
 - a. Move the parent TDE keystore from the original target host to the new host. If the vPDB is not being migrated to a new host, then this step is not needed. Similarly, if the new host has the parent keystore in a different location, then the parent keystore path for the vPDB needs to be updated.
 - b. Move the artifact directory for the vPDB from the original target host to the new host. If the vPDB is not being migrated to a new host, then this step is not needed.
 - c. Ensure that the new target container has the TDE keystore password set.
 - d. Merge the original target keystore into the new target keystore. This is required to support a rewind operation to a snapshot taken before the migrate.
4. Enable.

The diagram below illustrates the scenario of a migrated TDE-enabled vPDB.



Example

Consider the following TDE-enabled vPDB:

This vPDB is currently provisioned to the container database CDOMLOTG9620 on the host tde-target18. Connecting to that database, we can see the keys for the target container and the vPDB:

```
SQL> show pdbs
CON_ID  CON_NAME                OPEN MODE  RESTRICTED
-----
 2  PDB$SEED                READ ONLY  NO
 3  CDOMLOTG9620PDB1        READ WRITE NO
 4  CDOMLOTG9620PDB2        READ WRITE NO
 5  CDOMLOTG9620PDB3        READ WRITE NO
 6  TDE_VPDB                READ WRITE YES
SQL> select con_id, key_id from v$encryption_keys;
CON_ID  KEY_ID
-----
 1  ASFwmcfaMk8vv1LvzV0H8BEAAAAAAAAAAAAAAAAAAAAAAAAAAAA
 6  AdDdKibLKU9mv6PDAlVvH0AAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

The key for the target container starts with **ASFwmcfa** and the vPDB key starts with **AdDdKibL**. We also have an artifact directory for this vPDB found in `/work/oracle_tde_keystores/`:

```
$ ls /work/oracle_tde_keystores/
tde_vpdb_Encrypted_e9d2befb-9849-43c8-85f5-5ee8b760e334
```

We are going to migrate this vPDB to the container database CDOMLOTGCAE8 located on the host tde-target18b. Currently, that container has the following configuration:

```
SQL> show pdbs
CON_ID  CON_NAME                OPEN MODE  RESTRICTED
-----
```


```

2 PDB$SEED          READ ONLY  NO
3 CDOMLOTGCAE8PDB1  READ WRITE NO
4 CDOMLOTGCAE8PDB2  READ WRITE NO
5 CDOMLOTGCAE8PDB3  READ WRITE NO
SQL> select con_id, key_id from v$encryption_keys order by con_id;
CON_ID KEY_ID
-----
1 Ad344VSUDk/jv1VDb1QNBHMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
3 AetaL+IKpE+ov5avXouApwUAAAAAAAAAAAAAAAAAAAAAAAAAAAA
4 AXZTgaBbxk+6v6x2yHQ4ArgAAAAAAAAAAAAAAAAAAAAAAAAAAAA
5 AayTgbjkJE/3v82/hBBBAWEAAAAAAAAAAAAAAAAAAAAAAAAAAAA

```

The new target container has a CDB key starting with **Ad344VSU**. The migration steps are as follows:

1. Disable the vPDB in the original container.
2. Migrate the vPDB to point to the new container on the new host. The UI will report the following warnings during the migrate operation, which indicate the manual steps which need to be done:



3. Set the TDE keystore password for the new target container CDOMLOTGCAE8.
4. Move the parent keystore to the new target host tde-target18b. In this case, we will keep the existing location /u01/app/oracle/product/18.11.0.0/dbhome_1/admin/CDOMSHSR52CA/wallet on both hosts, so the configuration does not need to be updated.
5. Move the TDE artifact directory /work/oracle_tde_keystores/tde_vpdb_Encrypted_e9d2befb-9849-43c8-85f5-5ee8b760e334 to the new target host. If there are no TDE-enabled vPDBs on the new host, the oracle_tde_keystores directory may need to be created first.
6. Merge the original target keystore into the new target keystore. Oracle provides the ADMINISTER KEY MANAGEMENT MERGE KEYSTORE command to facilitate this. Note that you cannot successfully merge into a keystore that is currently in use by the database, so the recommended process is to first copy the original target keystore to the new host, merge the two keystores into a new keystore, replace the existing new target keystore and finally bounce the CDB to start using the merged keystore:
 1. Copy the original keystore to /tmp/tde-target18a on tde-target18b
 2. While connected to the new container, issue the merge keystore command:
 SQL> administer key management merge keystore '/tmp/tde-target18a' identified by *** and keystore '/u01/app/oracle/product/18.11.0.0/dbhome_1/admin/CDOMLOTGCAE8/wallet' identified by *** into new keystore '/tmp/merged' identified by ***;
 3. Copy the new merged keystore into the existing keystore location (backing up the existing keystore first).
 4. If there is an autologin wallet configured for the container, it must be recreated.
 5. Shutdown and restart the new target container database.
 - Confirm that the keystore now contains both the original key and the new key:

```
SQL> select key_id from v$encryption_keys where con_id = 1;
KEY_ID
-----
ASFwmcfaMk8vv1LvzV0H8BEAAAAAAAAAAAAAAAAAAAAAAAAAAAA
Ad344VSUDk/jv1VDb1QNBHMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
```

Note that we have the key starting with **ASFwmcfa** from the original container, and the key starting with **Ad344VSU** from the new container.

- Enable the vPDB to complete the migration operation. We can now see that the vPDB is successfully started in the new container, with its key starting with **AdDdKibL**:

```
SQL> show pdbs
CON_ID CON_NAME          OPEN MODE  RESTRICTED
-----
  2 PDB$SEED             READ ONLY NO
  3 CDOMLOTGCAE8PDB1    READ WRITE NO
  4 CDOMLOTGCAE8PDB2    READ WRITE NO
  5 CDOMLOTGCAE8PDB3    READ WRITE NO
  6 TDE_VPDB            READ WRITE NO
SQL> select con_id, key_id from v$encryption_keys order by con_id;
CON_ID KEY_ID
-----
  1 Ad344VSUDk/jv1VDb1QNBHMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  1 ASFwmcfaMk8vv1LvzV0H8BEAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  3 AetaL+IKpE+ov5avXouApwUAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  4 AXZTgaBbxk+6v6x2yHQ4ArgAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  5 AayTgbjkJE//v82/hBBBAWEAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  6 AdDdKibLKU9mv6PDAlVvVH0AAAAAAAAAAAAAAAAAAAAAAAAAAAA
6 rows selected.
```

Related Topics

- [Transparent Data Encryption and Delphix](#)
- [Provisioning a TDE-enabled vPDB](#)
- [Provisioning a vPDB from a Non-Multitenant Source](#)
- [Refreshing and Rewinding a TDE-enabled vPDB](#)

Migrating a vPDB (from a linked CDB) to a higher Oracle version (linked CDB)

Given a vPDB in a linked CDB (CDBold for hereon) the following steps can be used to migrate it to a higher Oracle version linked CDB (CDBnew from hereon):

- It might be necessary to run pre-upgrade Oracle scripts before doing the upgrade, consult Oracle documentation.
- Disable the vPDB, using the GUI or the CLI.
- Using the CLI perform manual migration. If the CDBnew is in the same host as the CDBold, then only cdbConfig needs to be specified, otherwise the repository and the environmentUser also need to be specified:

```
> cd /sourceconfig
sourceconfig> select vpdb
sourceconfig 'vpdb'> update
sourceconfig 'vpdb' update *> set cdbConfig=CDBnew
sourceconfig 'vpdb' update *> set environmentUser='<new oracle user>'
sourceconfig 'vpdb' update *> set repository='<new oracle home>'
sourceconfig 'vpdb' update *> commit
```

⚠ After enabling the vPDB without starting, the environment of the linked CDB should not be refreshed until the vPDB has been created and upgraded.

- Enable the vPDB without starting it:

```
> cd /source
source> select vpdb
source 'vpdb'> enable
source 'vpdb' enable *> set attemptStart=false
source 'vpdb' enable *> commit
```

- Add the vPDB to CDBnew using sqlplus, which has two options:
 - The vPDB can be plugged as a clone (to get a new DBID/GUID):

```
SQL> create pluggable database vpdb as clone using '/mnt/provision/vpdb-
CDOMLOSRA92E/datafile/delphix_group_writable/VPDB.xml' nocopy tempfile
reuse;
```

- The vPDB can be plugged without the clone option (to retain the same DBID/GUID):

```
SQL> create pluggable database vpdb using '/mnt/provision/vpdb-
CDOMLOSRA92E/datafile/delphix_group_writable/VPDB.xml' nocopy tempfile
reuse;
```

- Open the vPDB in upgrade mode:

```
SQL> alter pluggable database vpdb open upgrade;
```

- Upgrade the vPDB following the procedures documented by Oracle.
- Take a new snapshot with the new version upgraded and delete the previous snapshot containing the previous version metadata.

All Delphix operations on the upgraded vPDB are possible, which include:

- Start
- Stop
- Enable
- Disable
- Taking a snapshot
- Rewind (to an upgraded snapshot)
- Refresh (requires additional steps)
 - Disable the vPDB
 - Using the CLI to migrate the vPDB back to CDBold:

```
> cd /sourceconfig
sourceconfig> select vpdb
sourceconfig 'vpdb'> update
sourceconfig 'vpdb' update *> set cdbConfig=CDBold
sourceconfig 'vpdb' update *> set environmentUser=<old oracleuser>'
sourceconfig 'vpdb' update *> set repository='<old oracle home>'
sourceconfig 'vpdb' update *> commit
```

- Refresh the vPDB
- Repeat the procedure for migration as described before

V2P with an Oracle VDB

This topic describes the procedure for exporting a virtual database (VDB) to a physical one, also known as V2P.

Procedure

1. Log in to the **Delphix Management** application.
2. Click **Manage** and select **Datasets**.
3. Select the dSource or VDB you want to export.
4. Select the snapshot of the dSource or VDB state you want to export.
5. If you want to export the state of the database from a specific point in time, select the **LogSync** icon and then select the point in time from which you want to create the export.



6. From the actions menu (...) select **Virtual to Physical**.
7. Select the **target environment**. The target environment should have been added to Delphix previously and should meet all target host requirements, see [Requirements for Oracle Hosts and Databases](#).
8. Enter the Target Directory for the export.
The target directory you enter must exist in the target environment, and the Delphix operating system user listed under the environment must have permission to write to it. The target directory should be empty.
9. Select whether or not to **Open Database After Recovery**.
If you do not select this option, the Oracle database will not undergo open resetlogs, and the database will not be available for read/write access. This can be useful if the files are to be used to restore an existing data file for recovery purposes. You can use the scripts that are created in the target environment to complete the database open process at a later time. For more information, see [Manually Recovering a Database after V2P](#).
10. Click **Advanced** to customize **data transfer settings**, customize the **target directory layout**, enter any **database configuration** parameters or enter **file mappings** from the source environment to the target. For more information, see [Customizing Target Directory Structure for Database Export](#), [Configuration Settings for Oracle Virtual Databases](#) and [Customizing VDB File Mappings](#). The data transfer settings are described below:
 - **Compression** – Enable compression of data sent over the network. Default is **Off**.
 - **Encryption** – Enable encryption of data sent over the network. Default is **Off**.
 - **Bandwidth Limit** – Select the network bandwidth limit in units of megabytes per second (MB/s) between the Delphix Engine and the target environment. The default is **0**, which means no bandwidth limit is enforced.
 - **Number of Connections** – Select the number of transmission control protocol (TCP) connections to use between the Delphix Engine and the target environment. Multiple connections may improve network throughput, especially over long-latency and highly-congested networks. The default is **1**.
 - **Number of Files to Stream Concurrently** – Select the number of files that V2P should stream concurrently from the Delphix Engine to the target environment. The default is **3**.
11. Click **Next**.
12. Select whether you want to have an email sent to you when the export process completes.
13. Click **Submit**.

Post-Requisites

If you did not select **Open Database After Recovery**, follow the instructions in [Manually Recovering a Database after V2P](#) to complete the database open process.


Resumable V2P

Resumable V2P is a capability that allows you to suspend a V2P operation and then resume it at a later time, without redoing any of the work already completed. For example, any portion of a file that has already been

transferred to the target environment is not re-sent. For an entire file that has already been transferred, no part is re-sent.

The image below presents a progress bar, a stop button, and a pause button while a V2P is running. To manually suspend a V2P operation:

1. Click the **pause** button.



The next image represents a message alert generated after a V2P job has been suspended. To manually resume the job:

1. Click the **play** button.



Recoverable Errors

Broadly speaking, a "recoverable error" is an error condition caused by a disruption in the environment or on the target host, not errors in the actual V2P operation. Examples of recoverable errors include:

- A timeout due to a network outage
- Running out of disk space on the target environment
- An inability to create directories or files on the target environment

You can often address recoverable errors by taking some action to fix the problem, such as freeing up space on the target environment.

Auto-Suspend

A V2P operation that encounters a recoverable error is auto-suspended: it appears as a suspended job in the user interface (UI), with a message detailing the error condition. Once you have fixed the error, you can simply resume

the job. Alternatively, you can cancel the job. Just as when you manually suspend and resume a job, any portion of a file that has been transferred, including possibly the entire file itself, is not re-sent when the job resumes.

Related Topics

- [Customizing Target Directory Structure for Database Export](#)
- [Manually Recovering a Database after V2P](#)
- [V2P with Unstructured Files](#)
- [Requirements for Oracle Hosts and Databases](#)
- [Customizing Oracle VDB Environment Variables](#)
- [Customizing VDB File Mappings](#)

Move an Oracle VDB to a Physical ASM or Exadata Database

Overview

This article describes how to move a virtual database (VDB) into a physical database stored on Oracle Automatic Storage Management (ASM) disk groups. This is a scripted procedure that is assisted (but not fully automated) by Delphix. No intermediate storage is needed; the database files are moved directly from Delphix into the ASM diskgroup(s).

This procedure can be used to export a VDB to ASM disk group(s), including disk group(s) residing in an Oracle Exadata machine. The `move-to-asm.sh` script follows Oracle's recommended best practice of using a single disk group for data files. A separate disk group can be specified for redo log files. Professional Services can be engaged if customization is needed.

 This procedure applies to **all** Oracle RDBMS Versions supported by Delphix; in addition, it is not supported by PDB/Multi-tenant type databases.

Procedure

1. Create an ASM disk group that will contain all the database files. Optionally, create a separate disk group for redo log files.
2. Provision a VDB on the target machine that is running Oracle ASM or Exadata in order to do the export; this VDB will be deleted after the export is complete. This temporary VDB can be provisioned from an existing VDB. You can also provision it from a dSource, which allows you to achieve a full restore of the original source ASM database.
3. Login to the target machine (or a node of the target RAC cluster) where the VDB instance and ASM instance are running. The directory will be in the form:
`<toolkit_path>/Delphix_<engine_id><user><host|cluster>/databases/oracle/
 <db_unique_name>/<database_sid>/`

For a RAC database, be sure to use the directory path containing the SID of the instance running on the node where you have logged in.

4. Ensure that Oracle environment variables ORACLE_HOME, ORACLE_SID, and CRS_HOME (RAC only) are correctly set for the VDB that needs to be moved.
5. Execute the script `move-to-asm.sh` as the **Environment User** who provisioned the single instance VDB. For a RAC VDB, the **Environment User** selected to execute the `move-to-asm.sh` must be the Oracle installation owner. This is due to an Oracle restriction that only the installation owner can invoke `srvctl` to add or remove database configurations.

Alternatively, enter `move-to-asm.sh` as a **Post-Provision Hook Script** when provisioning the VDB. This will provision and move the VDB into ASM disk groups in a single flow. You must specify the `-noask` option to execute in non-interactive mode. For more information about hook scripts, visit the [Hook Scripts for Automation and Customization](#) article.

Command syntax: `move-to-asm.sh [-noask] [-parallel #] [-dbunique db_unique_name]
 <data_diskgroup> [<redo_diskgroup>]`

Arguments	Description
-----------	-------------

<code>-noask</code> [optional]	Do not prompt for confirmation before moving the VDB. The default is to prompt.
<code>-parallel</code> # [optional]	The number of RMAN channels used to move the VDB to ASM. The default is 8.
<code>-dbunique</code> <db_unique_name> [optional]	Unique database name for the resulting physical database. The default is the unique name of the VDB. The <code>-dbunique</code> argument only works if a database with that name that you select already exists.
<data_diskgroup> p> [required]	Target ASM disk group for data, server parameter and control files.
<redo_diskgroup> p> [optional]	Target ASM disk group for redo log files. Default is <data_diskgroup> .

move-to-asm script

The *move-to-asm.sh* script generates several output files. These files are all written to the working directory of the script:

`init<oracle_sid>_run<process-id>_moveasm.ora`

The log file for the operation:

`move-to-asm.sh_<oracle_sid>_run<process-id>.log`

The init.ora parameter file created for the ASM DB instance:

`source_init<oracle_sid>.ora`

The init.ora for the source database (from which the VDB was created).

Finishing Steps

The final steps to manually execute are shown when the script completes and included in the script output.

1. For Single Instance only: copy generated `init<$ORACLE_SID>.ora` parameter file to the default `$ORACLE_HOME/dbs/init<$ORACLE_SID>.ora`
2. Modify initialization parameters to match the original source database parameters, if necessary. As a convenience to assist with this step, the source database parameters are restored as `source_init<$ORACLE_SID>.ora` in the execution directory. Be sure to copy this file to a save place *before* deleting the temporary VDB.

The *move-to-asm.sh* script generates several output files. These files are all written to the working directory of the script and should also be copied to a safe place if needed:

`move-to-asm.sh_<oracle_sid>_run<process-id>.log`

The log file for the operation:

```
init<oracle_sid>_run<process-id>_moveasm.ora
```

The init.ora parameter file created for the ASM DB instance:

```
source_init<oracle_sid>.ora
```

The init.ora for the source database (from which the VDB was created).

3. Delete the Delphix VDB that was moved.
4. Start up the physical database that will now run on ASM.
(A RAC database is automatically started up by the `move-to-asm.sh` script using `srvctl` but may have been stopped when the VDB was deleted, if it has the same name.)

Sample Execution

```
$ cd toolkit/Delphix_564dc816_d457_6224_1327_4d43fd1b3d97_oracle_cluster/databases/
oracle/RAC2ASM/RAC2ASM1
$ export ORACLE_HOME=/u01/app/oracle/product/12.2.0.1/dbhome_1
$ export ORACLE_SID=RAC2ASM1
$ sh move-to-asm.sh -noask +DATA +REDO
END_OF_SETUP
=====
Virtual-to-ASM script (move-to-asm.sh v5.2) for Delphix
Copyright (c) 2013, 2018 by Delphix. All rights reserved.
=====

Moving database RAC2ASM to ASM: started at Wed Feb 19 16:08:43 EST 2020
db_unique_name => RAC2ASM
ORACLE_SID => RAC2ASM1
ORACLE_HOME => /u01/app/oracle/product/12.2.0.1/dbhome_1
Datafile diskgroup => +DATA
Online redo diskgroup => +REDO
RMAN Channels => 8

Create cluster database RAC2ASM
Generate script to move tempfiles to ASM
Generate script to drop old tempfiles
Generate script to drop offline tablespaces
Generate script to make read-only tablespaces read-write
Generate script to make tablespaces read-only again
Make read-only tablespaces read-write
Remove offline tablespaces
Updating server parameter file with ASM locations
Creating temporary initialization parameter file
Move spfile to ASM
Move datafiles to ASM: started at Wed Feb 19 16:09:41 EST 2020
Move datafiles to ASM: completed at Wed Feb 19 16:10:31 EST 2020
Stop cluster database RAC2ASM
Restart cluster database RAC2ASM
Remove old tempfiles
Move tempfiles into ASM
Move Online logs
Restore any read-only tablespaces
Shutdown database
```

Starting up RAC database RAC2ASM...

Database RAC2ASM moved to ASM: completed at Wed Feb 19 16:12:45 EST 2020

Final Steps to complete the move to ASM:

Source parameters are restored at /home/oracle/toolkit/

Delphix_564dc816_d457_6224_1327_4d43fd1b3d97_oracle_cluster/databases/oracle/RAC2ASM/RAC2ASM1/source_initRAC2ASM1.ora

```
*.audit_file_dest='/u01/app/oracle/admin/DBOMSRD914EA/adump'
*.audit_trail='DB'
*.cluster_database=TRUE
*.compatible='12.2.0'
*.control_files='+DATA/DBOMSRD914EA/CONTROLFILE/current.297.1027776391','+REDO/DBOMSRD914EA/CONTROLFILE/current.264.1027776391'
*.db_block_size=8192
*.db_create_file_dest='+DATA'
*.db_create_online_log_dest_1='+DATA'
*.db_create_online_log_dest_2='+REDO'
*.db_name='DBOMSRD9'
*.db_recovery_file_dest='+REDO'
*.db_recovery_file_dest_size=6989807616
*.db_unique_name='DBOMSRD914EA'
*.diagnostic_dest='/u01/app/oracle'
*.dispatchers='(PROTOCOL=TCP) (SERVICE=DBOMSRD914EAXDB)'
*.dlpx_recovery_log_archive_format='arch_%t_%s_%r.log'
*.instance_number=1
*.listener_networks=''
*.local_listener='(ADDRESS=(PROTOCOL=TCP)(HOST=10.43.5.243)(PORT=1521))'
*.log_archive_format='%t_%s_%r.dbf'
*.memory_max_target=1073741824
*.memory_target=1073741824
*.nls_language='AMERICAN'
*.nls_territory='AMERICA'
*.open_cursors=300
*.processes=300
*.remote_listener='node1.sample.delphix.com:1521'
*.remote_login_passwordfile='EXCLUSIVE'
*.thread=1
*.undo_tablespace='UNDOTBS1'
```

Copy **this** file to a safe location before removing the VDB

- 1) Delete VDB on Delphix.
- 2) Modify initialization parameters to match source and restart.

Related Topics

- [Provisioning and Managing Virtual Databases with Oracle](#)

Provisioning from a Replicated Oracle VDB

This topic describes how to provision from a replicated dSource or virtual database (VDB).

The process for provisioning from replicated objects is the same as the typical VDB provisioning process, except that first you need to select the replica namespace containing the replicated object.

Prerequisites

- You must have a dSource or VDB that has been replicated from one Delphix Engine to another, as described in [Replication Overview](#).
- The Delphix Engine containing the replicated dSource or VDB must have a compatible target environment that it can use to provision a VDB from the replicated dSource or VDB.

Procedure

1. On the Delphix Engine containing the replicated dSource or VDB, login to the **Delphix Management** application.
2. In the top menu bar, click **Manage**.
3. Select **Datasets**.
4. From the list of replica namespaces, select the **replica namespace** that contains the dSource or VDB from which you want to provision.
5. The provisioning process is now identical to the process for provisioning standard objects.

Post-Requisites

Once the provisioning job has started, the user interface will automatically display the new VDB in the live system.

Related Topics

- [Replication Overview](#)

Provisioning a TDE-enabled vPDB to a Cluster Target

Overview

Provisioning a Virtual Pluggable Database (vPDB) to a linked RAC container database first involves using the GUI or CLI to specify the vPDB parameters (such as the vPDB name and target container) along with the snapshot to provision from. Once the provision job is started with these parameters, the Delphix Engine does the following:

1. Chooses an available instance from the target cluster and mounts the snapshot files on that instance.
2. Creates and opens (in mount mode) the auxiliary container database on the target instance, using the snapshot files. The auxiliary container database will have both the CDB and PDB data files from the dSource.
3. Completes recovery to bring the auxiliary container database into a consistent state.
4. Finalizes the state of the auxiliary database and unplugs the vPDB datafiles.
5. Plugs the vPDB into the target database, and opens it in read-write mode on the same target instance used for recovery.
6. Spawns start jobs to open the vPDB in read-write mode on the remaining target instances in parallel.

If the dSource is TDE-enabled, then Delphix will need to perform additional operations to complete the provision of a TDE-enabled vPDB to a TDE-enabled cluster target (indicated in **red**):

1. Chooses an available instance from the target cluster and mounts the snapshot files on that instance.
2. **Creates a keystore with the necessary keys to apply encrypted archived log files on the target instance.**
3. Creates and opens (in mount mode) the auxiliary container database on the target instance, using the snapshot files. The auxiliary container database will have both the CDB and PDB data files from the dSource.
4. Completes recovery to bring the auxiliary container database into a consistent state.
5. **Rotates the vPDB master encryption key by generating a new key that is unique to the vPDB and not associated with the source PDB.**
6. **Exports only the newly generated vPDB key to an exported keyfile to enable unplug.**
7. Finalizes the state of the auxiliary database and unplugs the vPDB datafiles.
8. **Imports the vPDB key from the exported keyfile into the target keystore.**
9. Plugs the vPDB into the target database, and opens it in read-write mode on the same target instance used for recovery.
10. Spawns start jobs to open the vPDB in read-write mode on the remaining target instances in parallel.

All the same information needed for a single instance TDE-enabled vPDB provision is also required for a cluster TDE-enabled vPDB provision, specifically the target keystore password, parent keystore path and password, and encryption secret. The keystores root path is required for a cluster provision.

Shared Storage for Keystores

In a cluster database, the database files are on shared storage, which is accessible from all instances in the cluster. If the database is encrypted, then the keystore file itself is also located on the operating system. Oracle recommends that the keystore also be on shared storage, on a different disk from the database files. If the keystore is not on shared storage, then it must be copied to all instances in the cluster after any changes, such as importing a key or generating a new key. Similarly, Delphix recommends that the parent keystore specified for the provision also be on shared storage. If not, then the same file must be copied to all of the instances before the vPDB is first provisioned, and any updates to the parent keystore must also be copied to all of the instances before any vPDB refresh or rewind.

As the autologin wallet is located in the same location as the password-based keystore, it should also be on shared storage. For this reason, `local_autologin` wallets will not work properly, as they will be accessed from multiple nodes in the cluster.

Keystores Root Path Requirements

The artifact directory for a given vPDB is created under the keystores root. As a subsequent operation on the vPDB may choose a different instance than the one used for the initial provision, the artifact directory needs to be accessible from all instances in the cluster. Delphix requires that the keystores root be specified for a TDE-enabled vPDB provision to a cluster target, and furthermore that it be located on shared storage. The engine will validate that this is the case before proceeding with the initial provision.

Oracle Hook Operations

Oracle RAC

When linking from, or provisioning to Oracle RAC environments, hook operations will not run once on each node in the cluster. Instead, the Delphix Engine picks a node in the cluster at random and guarantees all operations within any single hook will execute serially on this node.

Note that the Delphix Engine does not guarantee the same node is chosen for the execution of every hook but does guarantee that Pre-/Post- hook pairs (such as Pre-Sync and Post-Sync) will execute on the same node.

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at `/bin/sh`. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh "$ARG_ENVIRONMENT_VARIABLE"
"second argument in double quotes" 'third argument in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a `bash` binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

Shell Operation Tips

Using `nohup`

You can use the `nohup` command and process backgrounding from resource in order to "detach" a process from the Delphix Engine. However, if you use `nohup` and process backgrounding, you MUST redirect `stdout` and `stderr`.

Unless you explicitly tell the shell to redirect `stdout` and `stderr` in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either `stdout` or `stderr`. Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your `RunCommand` operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The RunExpect operation executes an Expect script on a Unix environment. The Expect utility provides a scripting language that makes it easy to automate interactions with programs which normally can only be used interactively, such as `ssh`. The Delphix Engine includes a platform-independent implementation of a subset of the full Expect functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Management application and CLI to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an `ssh` session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

Oracle Environment Variables

to access the dSource or VDB.

dSource Environment Variables

Environment Variable	Description
<code>ORACLE_SID</code>	The SID of the dSource
<code>ORACLE_BASE</code>	The home directory of the Oracle software hosting the dSource
<code>ORACLE_HOME</code>	The Oracle Home for the dSource
<code>CRS_HOME</code> (only set for RAC dSources)	The home directory for cluster services hosting the dSource

Environment Variable	Description
ORAENV_ASK	Always set to NO
DELPHIX_DATABASE_NAME	The database name reported by Oracle
DELPHIX_DATABASE_UNIQUE_NAME	The database unique name reported by Oracle
DELPHIX_PDB_NAME (only set for PDBs)	The PDB name reported by Oracle

VDB Environment Variables

Environment Variable	Description
ORACLE_SID	The SID for the VDB
ORACLE_BASE	The home directory for the Oracle software hosting the VDB
ORACLE_HOME	The Oracle Home for the VDB
CRS_HOME (only set for RAC VDBs)	The home directory for cluster services hosting the RAC VDB
ORAENV_ASK	Always set to NO
DELPHIX_DATABASE_NAME	The database name reported by Oracle
DELPHIX_DATABASE_UNIQUE_NAME	The database unique name reported by Oracle
DELPHIX_PDB_NAME (only set for PDBs)	The PDB name reported by Oracle
DELPHIX_MOUNT_PATH	The root of the NFS mount hosting the VDB data
MASKING_CONNECTOR_HOST (only set for masked provisioning)	The host that DMSuite will use for the connector

Environment Variable	Description
MASKING_CONNECTOR_PORT (only set for masked provisioning)	The port that DMSuite will use for the connector

❗ PATH and LD_LIBRARY_PATH configuration

`PATH` is configured by appending the `bin` directory in the Oracle home for the dSource or VDB.

`LD_LIBRARY_PATH` is configured by appending the `lib` directory in the Oracle home for the dSource or VDB.

Related Topics

- [Hook Scripts for Automation and Customization](#)

SQL Server Environments and Data Sources

This section covers the following topics:

- [Delphix Architecture with SQL Server](#)
- [Quick Start Guide for SQL Server \(Microsoft SQL Server on Windows\)](#)
- [SQL Server Support and Requirements](#)
- [Managing SQL Server Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with SQL Server](#)
- [Provisioning and Managing Virtual Databases with SQL Server](#)
- [Hooks for SQL Server](#)
- [Using HostChecker to Validate Target Database Servers](#)

Delphix Architecture with SQL Server

SQL Server is Microsoft's relational database which is typically run on Windows Server hosts. With Delphix, you can use various SQL Server configurations, ranging from Failover Clusters to Availability Groups. In this section, you'll find an overview of how Delphix works with SQL Server.

There are three key concepts when using Delphix with any data platform:

1. Environments: The server and software required to run a data set.
1. Source Environment: Source data to be ingested into Delphix. These will be used to create dSources.
2. Target Environment: Target hosts to provision VDBs.
 - dSources: A database that the Delphix Virtualization Engine uses to create and update or maintain virtual copies of your database
 - VDBs: A database provisioned from either a dSource or another VDB which is a copy of the source data. A VDB is created and managed by the Delphix Virtualization Engine.

With these concepts in mind, let's explore how Delphix connects to SQL Server environments and creates SQL Server dSources and VDBs.

Staging Push Mechanism with SQL Server

Starting Delphix Engine version 6.0.13.0 a new data ingestion mechanism has been introduced that will help users to push data into the staging database on their own.

The previous data ingestion mechanism of the SQL Server has a few limitations like dependency on the source access and limited backup vendor's support. Currently, Delphix supports Native, Commvault, NetBackup, Lightspeed, and Redgate backup vendors.

With Staging Push implementation, we have removed the dependency of accessing the customer's production database and also enabled the customers who are using the backup vendors that Delphix does not support.

To summarize the Staging Push mechanism.

- Customers now have ownership of the staging databases.
- Customers are now responsible to keep the staging database in sync with the source database. Note that the database files of the staging database are stored on Delphix Storage.
- Delphix is still responsible to snapshot the underlying data files and gathering any metadata required for provisioning from the snapshots.

Related Topics

- [Delphix in Multi-domain Windows Environments](#)
- [SQL Server Support and Requirements](#)

Delphix in Multi-domain Windows Environments

General Overview

When considering the Delphix logical architecture, there are four primary components:

1. Source host(s)
2. Delphix Virtualization Engine
3. VDB Target host(s)
4. Staging Target Host(s)

In SQL Server environments, the staging target host is used for staging data from the source database on the source host into Delphix. Although you can use any VDB target host on which the Delphix Connector service has been installed for this purpose, Delphix recommends a dedicated Staging Target Host for load isolation and separation of roles.

This page focuses on the process of getting source SQL Server data into the storage of the Delphix Virtualization Engine (DVE) via the Staging Target Host.

When considering SQL Server deployments in different enterprise environments, we often see cases where the production, development, test, or reporting environments exist in different Windows domains which may or may not have trust relationships. Such varying domain approaches can come into play due to security, organizational, geographical, or other technical reasons, and can make communication between Windows hosts more complicated to manage. Delphix is flexible enough to work in many configurations, but we want to help you choose the solution that best suits your unique environment.

We listed the four primary components of the Delphix logical architecture for SQL Server above. In addition, a fifth component in the Delphix logical architecture might be considered for use-cases #3, #4, and #5 in the Technical Overview below: a Connector host. The function of the Delphix Connector on that host is the discovery of the source environment via remote registry and ODBC calls. There are no Delphix software installation requirements for Windows source hosts, but it might be helpful to note this role can co-reside directly on the Windows source host for consolidation purposes if desired.

Technical Overview

Keep in mind that the Delphix Engine is always syncing with backups of the source database. It is never the live data that is ingested; it is always backups of different flavors.

If SQL Server simple recovery mode is used, these can be full or differential backups initiated by the source database. If full recovery mode is enabled, the Delphix Engine will typically leverage only transaction log backups after the initial data load. Again, the source database would initiate backups, and the Delphix Engine would collect the backup files that have been created by SQL Server. This approach of using transaction logs minimizes spikes in system load by ingesting smaller backups more often. Another option is copy-only backups, which the Delphix Engine initiates in a configuration known as Delphix Managed Backups. For more information refer to [Delphix as a Backup Solution to SQL Server](#).

Delphix can ingest database and log data from native backups, as well as a number of third-party backup products. SQL Server restores the backups onto the shared Delphix storage on the staging target host running the databases in recovery mode. We call this process “validated sync,” which is why you may hear the staging target also referred to as a validated sync server.


It is important to note that the Delphix Virtualization Engine (based on DxOS, itself derived from a UNIX-based OS) is not a domain member itself. The credentials we discuss in this document are between Windows servers, and the key domain-specific authentication is between the staging host and the UNC path to the SMB share where the backup data is stored.

In the rest of this section, we will describe multiple scenarios. Review them to determine which will fit best in your environment.

Case 1: Staging Target in Test Environment

In this case, we will review an environment with two domains: PRODUCTION and TEST, which have a domain trust relationship. This is one of the simplest and most straightforward approaches, as illustrated in the *Staging Target in Test Environment diagram below.*

In this example, the staging target host exists in the non-production TEST domain, but because of domain trust, accounts located in that domain can access resources in the PRODUCTION domain. This would allow the staging target host to connect to the PRODUCTION source host both for environment discovery and to the shared backup location “\\source\backups” over Server Messaging Block (SMB) to access database and transaction log backups.




Staging Target in Test Environment

Case 2: Staging Target in Production Environment

The scenario illustrated in the *Staging Target in Production Environment diagram below* shows a TEST domain that does not have access to resources in the PRODUCTION domain. However, the customer has determined that Delphix VDBs must be provisioned to the TEST domain. In this scenario, you can use the staging target host in the PRODUCTION domain to link to the PRODUCTION database and perform the normal restore of the DB and/or log files to the Delphix storage. You can then provision VDBs in the TEST domain.

In this case, VDBs can be completely isolated from the PRODUCTION domain, and there is no requirement for hosts in the TEST domain to have any direct access to resources in the PRODUCTION domain.




Staging Target in Production Environment

Case 3: Domain-agnostic Storage

This example shows a shared backup location that is not dependent on trust relationships between the PRODUCTION and TEST domains. Because Delphix uses UNC paths, it can support any protocol which provides UNC access for that backup data access – for example, SMB or iSCSI.

This is shown by the *Domain-agnostic Storage* diagram below by the arrow – stretching from bottom-left toward the upper-right and crossing the domain boundary – representing any UNC-compatible protocol connecting the staging target host to the data on the NAS host. Provided that the Delphix environment users on both the source host and staging target host have read/write access to the shared backup location on network-attached storage (NAS), the SQL Server instance running on the staging target host will be able to access the backup files needed.

Although this option is not specific to this case, you may notice we separated a connector role to its own connector host. As you can infer from the diagram, the Delphix Connector’s primary function on that host is the discovery of the source environment via remote registry and ODBC calls. Despite the fact that there are no software installation requirements for the source hosts in PRODUCTION, it may be helpful to note that you can even install this role directly on the source server for consolidation if you want to.




Domain-agnostic Storage

Case 4: Migrating Backup Files

In this somewhat more complex configuration, backup files are sent to storage in the PRODUCTION domain, while the host used to link to the source and perform the validated sync is in an isolated TEST domain. We have used a separate connector host in the PRODUCTION domain again, to perform the environment discovery of the source host there. Backup files for SOURCE are being stored on NAS.

We will link using the Staging Target Host and create VDBs in the TEST domain. When the Delphix Engine discovers that a new backup of PRODUCTION has been taken, it will attempt to find the relevant files in the shared backup location provided during linking. It does this by periodically performing a recursive search for the file names on the shared backup location. If it does not find the specific files, it will try again later. Knowing this, we can specify a shared backup location in the TEST domain and set up an automated process to copy the backup files from \nas\backups in the PRODUCTION domain to \\staging\backups in the TEST domain. We can use any copy mechanism to transfer the files, such as FTP or ROBOCOPY. The files must be available long enough for the Delphix Engine to detect and apply them to the recovery database on the Staging Target Host before removal.

We have customers who also use this model in cases with multiple data centers (on-premise deployments) or virtual private clouds (cloud deployments) rather than multiple domains. These customers want database and transaction log backups to be available in secondary data centers or private clouds, but they want to make sure that the data is only copied over the WAN once.



Migrating Backup Files


Case 5: SMB Anonymous Access

In this example (shown by the *SMB Anonymous Access* diagram below), a Windows SMB connection is traversing domains that do not have a trust relationship. This approach is problematic because there is no simple configuration for SMB file sharing that does not rely on domain trusts. As a result, there is no way to specifically grant accounts in the TEST domain access to SMB shares in the PRODUCTION domain.

Because such users cannot be authenticated, they are treated as “anonymous” users and do not have permission to any resources by default.

Windows provides an “Everyone” group. However, this group still only applies to accounts that can be authenticated in the domain, so you cannot use that group in this case. There is still a way to configure access to the shared backup location on `\\source\backups` by accounts in other domains, such as TEST. However, because it relies on anonymous access, you will need to consider the security implications of enabling this method, as well as measures that could mitigate any additional risk in your environment – for example, a private VLAN or IPSEC between hosts.

1. Enable the “Guest” account on the server source Server – for example, `\\SQLPROD`.
2. Create a share where full and transaction log backups will be stored – for example, `\\SQLPROD\backups`.
3. Configure read-only security access for both the folder security permissions on the shared directory and the share permissions for the “guest” account.



SMB Anonymous Access

Here are some additional links from Microsoft that relate to anonymous sharing:

- [Network access: Let Everyone permissions apply to anonymous users](#)
- [Network access: Shares that can be accessed](#)

Related Topics

- [Delphix Architecture with SQL Server](#)

Quick Start Guide for SQL Server (Microsoft SQL Server on Windows)


This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with SQL Server database objects in the Delphix Engine. It does not cover any advanced configuration options or best practices, which can have a significant impact on performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix functionality. It assumes you will use the VMware Hypervisor. It assumes you are running supported combinations of software as explained here: [Supported OS, SQL Server, and Backup Software Versions for SQL Server](#).

Overview

In this guide, we will walk through deploying a Delphix Engine, starting with configuring Source, Staging (aka Validated Sync), and Target database environments on Windows servers. We will then create a dSource, and provision a VDB.

The following diagram describes the engine topology for SQL Server environments. It illustrates the recommended ports to be open from the engine to remote services, to the Delphix Engine, and to the Source, Target and Validated Sync Environments.

For purposes of the QuickStart, you can ignore any references to Replication or Masking, such as the engines shown in the diagram below.



Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online.
This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
    example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
    the primary interface. Setting this value
                    to "true" will cause all other properties (address,
hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
"1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
-- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24

```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
```

```
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

- Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, 192.168.1.2/24)

- Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

- Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the
browser at:
```

```
http://<address>
```

Type "exit" to disconnect, or any other commands to **continue** using the CLI.


- Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
- Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Once you set up the network access for Delphix Engine, navigate to the Delphix Engine URL in your browser for server setup.

The welcome screen below will appear for you to begin your Delphix Engine setup.



The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at <http://<Delphix Engine>/login/index.html#serverSetup>.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

⚠ The default domain user created on Delphix Engines from 5.3.1 is known as **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

System Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications.

Choose your option to set up system time in this section. For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click **Next** to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.
4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

⚠ Although your Delphix Engine will work without registration, we strongly recommend that you register each Delphix Engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

SQL Server Source Hosts and Databases

Source Host Requirements

Windows servers that will be added as Source Environments must meet the following requirements:

Source Host Requirement	Explanation
<p>To allow Delphix-initiated backups, the service account running each SQL Server instance (the Instance Owner) should be one of:</p> <ul style="list-style-type: none"> • A domain user (e.g. <code>MYDOMAIN\accountname</code>) (RECOMMENDED) • A Managed Service Account or Group Managed Service Account (<code>MYDOMAIN\accountname\$</code>) (requires Windows 2012 and later, and SQL Server 2008R2 or later) • The LOCAL SYSTEM account (<code>NT AUTHORITY\SYSTEM</code>) • The NETWORK SERVICE account (<code>NT AUTHORITY\NETWORK SERVICE</code>) 	<p>Backups initiated by the Delphix Engine (Delphix Managed Backups or manual snapshots which request a backup) will fail if the Delphix Engine cannot map the Instance Owner to an Active Directory user or computer object.</p>
<p>The source host, proxy, and staging environments must have appropriate cross-domain trust relationships</p>	<p>For more information on these requirements, see the document Delphix in Multi-domain Windows Environments.</p>

Source Windows User Requirements

⚠ Authenticate to Active Directory

Authenticate to Active Directory using a username and password: Delphix authenticates to Active Directory using a username and password. The Windows domain user cannot authenticate using Managed Service Accounts, Group Managed Service Accounts, smart cards, or other authentication methods. The Delphix Engine can retrieve passwords from a supported password vault, as described in [Password Vault Support](#).

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding source environments to the Delphix Engine. The Windows domain user provides remote read-only access to the Windows Registry, which allows the discovery of SQL Server instances and databases (note: Delphix also uses the registry to discover cluster nodes in AG Cluster Source Environments). This user must have the following permissions:

Source Windows User Requirement	Explanation
Be a member of the Backup Operators group on the Source Host	Delphix uses this privilege to discover SQL Server instances and gather system details, using Windows remote registry access.
Be a member of the Users group on the Staging Host	In order to discover and query SQL Server instances as the Source Windows User, scripts are run on that user at the Staging Host.
Have the "Log on as batch" privilege on the Staging host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Source Database Login Requirements</i> section below.

Source Database Login Requirements

The Delphix Engine requires SQL Server logins to be created on each SQL Server instance that the Delphix Engine will communicate with:

- A database login for Environment discovery and monitoring, specified when Adding an Environment. This must be a Windows Authentication login for the Source Windows User configured in the previous section.
- A database login for dSource (Source Database) monitoring and interaction, specified when Linking a dSource. This user can be:
 - The same as the Source Windows User;
 - A different Windows Authentication login (this user must also have Log on as a Batch Job privileges on the Staging host); or
 - A SQL Authentication login

These users must have the following permissions on each instance:

Object	Privileges Required	dSource User	Environment User	Purpose
Server	CONNECT SQL	✓	✓	Access to the SQL Server instance
Database: master	db_datareader	✓	✓	Access to information about attached databases
Database: msdb	db_datareader	✓		Access to the backup history
Each user database to be linked	PUBLIC	✓		Delphix will periodically run queries to check the current size of the database
Each user database to be linked	db_backupoperator	✓		Optional: Required for backups to be initiated by Delphix (using Delphix Managed Backups, or when Delphix initiates a backup during a manual Snapshot)
Server	VIEW ANY DEFINITION	✓	✓	Optional: Required for the discovery of databases in Availability Groups
Server	VIEW SERVER STATE	✓	✓	Optional: Required for the SnapSync and discovery of Availability Groups
Object: master.dbo.sqlutility	EXECUTE	✓		Optional: Required when using backups created by Red Gate SQL Backup
Object: master.dbo.xp_sqlightspeed_version	EXECUTE	✓		Optional: Required when using backups created by Quest LiteSpeed for SQL Server

List of Source tables accessed by the Delphix Engine

Using the db_datareader permission, the Delphix Engine accesses the following system tables in the master and msdb databases on the source host:

System table	Justification
master.sys.databases	Used to determine the name and recovery model of databases within discover SQL Server instances
master.sys.availability_groups	Used for discovering all the availability groups within an Availability Group source environment.
master.sys.availability_group_listeners	Used for discovering all the availability group listeners within an Availability Group source environment. <ul style="list-style-type: none"> A requirement for dSource linking of SQL Server clustered databases (replicas) is to provide an AG (Availability Group) listener for AG cluster source discovery. This is implemented on AG cluster source discovery as a failsafe if AG cluster source database authentication configuration change down the line, ensuring the Delphix engine has a way to reach the cluster and continue certain operations.
master.sys.availability_databases_cluster	Used for discovering all the availability group clusters within an Availability Group source environment.
master.sys.availability_replicas	Used for discovering all the availability group replicas within an Availability Group source environment.
master.sys.database_files	Used to determine the size of databases and whether filestream is enabled for a database
master.sys.dm_exec_requests	Used to enable Delphix to report backup operation progress
master.sys.master_files	Used to determine the primary file of a database
master.sys.filegroups	Used to determine the filegroups of a database so that Delphix can configure VDBs with the same filegroups
msdb.dbo.backupset	Used to determine new backups that have been taken. This table is regularly queried to find out if a new backup image has been taken and needs to be synchronized with Delphix.
msdb.dbo.backupmediafamily	Used to determine the physical device names of the backup files comprising a backup.

SQL Server Staging Hosts and Databases

Staging Host Requirements

Windows servers which will be added as Staging Environments must meet the following requirements:

Staging Host Requirement	Explanation
<p>The service account running each SQL Server instance (the Instance Owner) must be one of:</p> <ul style="list-style-type: none"> • A domain user (e.g. <code>MYDOMAIN\accountname</code>) (RECOMMENDED) • A Managed Service Account or Group Managed Service Account (<code>MYDOMAIN\accountname\$</code>) (requires Windows 2012 and later, and SQL Server 2008R2 or later) • The LOCAL SYSTEM account (<code>NT AUTHORITY\SYSTEM</code>) • The NETWORK SERVICE account (<code>NT AUTHORITY\NETWORK SERVICE</code>) 	<p>Access to existing database backups for Snapshots and Validated Sync operations will fail if the service account cannot access backups created by the Source database instance.</p>
<p>The source host, proxy and staging environments must have appropriate cross-domain trust relationships</p>	<p>For more information on these requirements, see the document Delphix in Multi-domain Windows Environments.</p>
<p>The edition of the installed SQL Server instance(s) must support all database features used by linked Source databases</p>	<p>SQL Server may raise an error during some dSource operations if the Staging SQL Server Instance does not support features used by the Source Database.</p> <p>This is most easily addressed by using the same edition of SQL Server as the Source database.</p> <p>Features in use on the Source database can be checked using the <code>sys.dm_db_persisted_sku_features</code> dynamic view.</p>
<p>Must have both the Source and Staging Windows Users configured as local Windows users</p>	<p>See the sections <i>Source Windows User Requirements</i> and <i>Staging Windows User Requirements</i> for more detail.</p>
<p>Delphix Connector software is installed and running</p>	<p>See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.</p>

Staging Host Requirement	Explanation
Recommended iSCSI Registry settings must be in place	See Requirements for Windows iSCSI Configuration and Knowledge Base Article KBA1251 for instructions on applying these settings*. *Please note that in 6.0.x, these performance settings are subject to change based on further tuning efforts by Delphix Engineering.
sqlcmd command line utility must be installed on the servers hosting VDBs and staging databases. The utility should be in the Delphix operating system user's PATH environment variable.	Delphix does not run sqlcmd (or any other process) directly on the source SQL Server instance. It runs sqlcmd on the staging/connector host.

Staging Windows User Requirements

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding Staging environments to the Delphix Engine. This user must have the following permissions:

Staging User Requirement	Explanation
Be a member of the Windows "Local Administrators" group on the Staging Host	We require this permission for mounting iSCSI LUNs presented by the Delphix Engine to the staging and target hosts. Microsoft utilities used by the Delphix Engine, such as diskpart, require membership of this group .
Have the "Log on as batch" privilege on the Staging host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to access existing backups from the Source Database	If accessing existing native, Red Gate, or LiteSpeed backups, the following permissions are required: <ul style="list-style-type: none"> The Staging Windows User and the service account running SQL Server must both have permission to access the database backups via SMB (Windows file sharing) The Staging Windows User and the service account running SQL Server must both have NTFS Permissions to access the database backups Separate documents describe requirements if Linking a dSource from a NetBackup SQL Server Backup or Linking a dSource from a Commvault SQL Server Backup .
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Staging Database Login Requirements</i> section below.

Staging Database Login Requirements

The Delphix Engine requires a Windows Authentication login to be created for the Staging Windows User on each SQL Server instance that the Delphix Engine will communicate with, with the following permissions:

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	CONNECT SQL	✓	Access to the SQL Server instance
Server	sysadmin	✓	<p>The staging and target databases are managed and administered completely by Delphix. Our functionality requires many administrative operations on those databases and requires full access to them.</p> <p>Since database ownership can be changed by customers as part of configuring virtual databases, we require the sysadmin role to continue to administer the databases.</p>

SQL Server Target Hosts and Databases

Target Host Requirements

Windows servers which will be added as Target Environments must meet the following requirements:

Target Host Requirement	Explanation
Delphix Connector software is installed and running	See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.
Recommended iSCSI Registry settings must be in place	See Requirements for Windows iSCSI Configuration and Knowledge Base Article KBA1251 for instructions on applying these settings.
The edition of the installed SQL Server instance(s) must support all database features used by linked Source databases	<p>SQL Server may raise an error during some dSource operations if the Staging SQL Server Instance does not support features used by the Source Database.</p> <p>This is most easily addressed by using the same edition of SQL Server as the Source database.</p> <p>Features in use on the Source database can be checked using the <code>sys.dm_db_persisted_sku_features</code> dynamic view.</p>

Target Host Requirement	Explanation
sqlcmd command line utility must be installed on the servers hosting VDBs and staging databases. The utility should be in the Delphix operating system user's PATH environment variable.	Delphix does not run sqlcmd (or any other process) directly on the source SQL Server instance. It runs sqlcmd on the staging/connector host.

Target Windows User Requirements

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding Target environments to the Delphix Engine. This user must have the following permissions:

Target User Requirement	Explanation
Be a member of the Windows "Local Administrators" group on the Target Host	We require this permission for mounting iSCSI LUNs presented by the Delphix Engine to the staging and target hosts. Microsoft utilities used by the Delphix Engine, such as diskpart, require membership of this group .
Have the "Log on as batch" privilege on the Target Host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Target Database Login Requirements</i> section below.
Delphix Connector software is installed and running	See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.

Target Database Login Requirements

The Delphix Engine requires a Windows Authentication login to be created for the Target Windows User on each SQL Server instance that the Delphix Engine will communicate with, with the following permissions:

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	CONNECT SQL	✓	Access to the SQL Server instance

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	sysadmin	✓	<p>The staging and target databases are managed and administered completely by Delphix. Our functionality requires many administrative operations on those databases and requires full access to them.</p> <p>Since database ownership can be changed by customers as part of configuring virtual databases, we require the sysadmin role to continue to administer the databases.</p>

Supported Roles for Failover Cluster Instances and Always On Availability Groups

Failover Cluster Instances and Always On Availability groups cannot be used as Staging Environments, and VDBs cannot be provisioned into Availability Groups.

 When adding a Failover Cluster Instance or Always On Availability Group, all nodes of the cluster must meet the requirements described in this document.

The following table shows how different SQL Server instance types should be added from the Add Environment screen:

Color		Supported?		
✓		Yes		
✗		No		
		Environment Role		
Instance Type	Added As	Source Environment	Staging Environment	Target Environment
Standalone Instance	Standalone	✓ *	✓	✓
Failover Cluster Instance (FCI)	Standalone	✓ *	✗	✗
Failover Cluster Instance (FCI)	Cluster	✗	✗	✓

		Environment Role		
Always On Availability Group (AG)	Cluster	Supported		

* Databases that are participating in Availability Groups will not be discovered during the discovery of a Standalone environment.

** VDBs cannot be provisioned into availability groups. However, SQL Server instances that participate in an Availability Group can also be added as Standalone Instances.

Using a Failover Cluster Instance as both Source and Target

A Failover Cluster Instance added as an environment once (as either a Source or Target environment) cannot be used as both a Source and Target.

If this is required, the environment can be added twice:

- Once as a Standalone Source environment
- Once as a Cluster Target environment

The Standalone Source environment can be used for linking dSources, and the Cluster Target environment is used for provisioning VDBs.

As suggested by the Best Practice note earlier in this article, this is not a recommended configuration.

Where possible, SQL Server failover cluster instances that the Delphix Engine will use as a target should not be used to host databases other than Delphix VDBs.

Requirements for Failover Cluster Target Environments

The following additional requirements exist for Windows Failover Cluster Instances added as Target Environments, to be used for VDB Provisioning.

- You must first add each node in the Window Failover Cluster individually as a standalone target environment, using a non-clustered address. See [Adding a SQL Server Standalone Target Environment](#).
 - The Delphix Engine may show a warning that it cannot discover the Failover Cluster Instances on each standalone host. This is expected.
- Each clustered SQL Server instance must have at least one clustered disk added to the clustered instance resource group, which can be used for creating mount points to Delphix storage.
 - The clustered drive must have a drive letter assigned to it.
 - The clustered drive must be formatted using the "GUID Partition Table (GPT)" partition style, in order for the Delphix Engine to automatically discover the drive letter as a valid option for the cluster instance. An MBR-formatted disk requires manual verification outside of Delphix that the disk has been correctly added to the MSSQL clustered resource group prior to creating the VDB. When provisioning the VDB, you must manually specify the desired MBR disk, because it will not appear in the Delphix GUI.
 - The clustered drive must be added to the clustered instance resource group as a dependency in the Failover Cluster Manager.
- Each node in the cluster must have the Failover Cluster Module for Windows PowerShell feature installed.
 - While running Windows PowerShell as an administrator, enter this command to test that the module is available: `Import-Module FailoverClusters`
- An additional target environment that can be used as a Connector Environment must exist.
 - This environment must **not** be a node in the cluster.
 - This environment should be part of the same Active Directory domain as the cluster.

Add the SQL Server Source Environment

Delphix does not require running the Connector on your source. Instead, you'll use the Validated Sync environment as a connector environment to discover the source by proxy.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions** menu, and select **Add Environment**.
5. In the **Add Environment** wizard, Host and Server tab select:
 - a. Host OS: **Windows**
 - b. Host Type: **Source**
 - c. Server Type:
 - If you are adding a Windows Server Failover Cluster (WSFC), add the environment based on which WSFC feature the source databases use:
 - Failover Cluster Instances
Add the environment as a **standalone** source using the **cluster name** or **address**.
 - AlwaysOn Availability Groups
Add the environment as a **cluster** source using the **cluster name** or **address**.
 - Otherwise, add the environment as a **standalone** source.
6. Click **Next**.
7. In the Environment Settings tab select a **Connector Environment**.
Connector environments are used as a proxy for running discovery on the source. If no connector environments are available for selection, you will need to set them up as described in [Adding a SQL Server Standalone Target Environment](#). Connector environments must:
 - have the Delphix Connector installed
 - be registered with the Delphix Engine from the host machine where they are located.
8. Enter the **Environment Name**, **Node Address**, **OS Username**, and **OS Password** for the source environment.
9. Click **Submit**.

As the new environment is added, you will see multiple jobs running in the Delphix Admin Job History to Create and Discover an environment. In addition, if you are adding a cluster environment, you will see jobs to Create and Discover each node in the cluster and their corresponding hosts. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel. If you don't see it, click the **Actions** menu and select **Refresh All**.

Linking a SQL Server Data Source (dSource)

Linking a dSource will ingest data from the source and create a dSource object on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

For an overview of all dSource related actions, please [Managing Data Sources and Syncing Data](#).

When linking a dSource from a SQL Server source database, Delphix offers several different methods of capturing backup information:

- SQL Server Managed Backups, where the SQL Server source database schedules and initiates backups and the Delphix Engine captures them
- Full backups
- Full or differential backups
- Transaction log backups (with LogSync disabled)
- Transaction log backups (with LogSync enabled)

- Delphix Managed Backups, where the Delphix Engine schedules and initiates the backups from the source database, and captures them

Procedure

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.
4. In the **Add dSource** wizard, select the source database with the correct environment user-specified.
5. Select user type for source database authentication and enter the login credentials. Enter username and password for Database user or Domain (Windows) user. For Environment User, select a source environment user from the dropdown list and click **Next**.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed. For more information, [Data Management Settings for SQL Server Data Sources](#).
8. Select the Staging environment and SQL Instance that will be used to manage the staging database used for validated sync of the dSource.
9. Select any policies for the new dSource.
10. Enter any scripts that should be run on the **Hooks** page.
11. Review the dSource Configuration and Data Management information, and then click **Submit**.

Provisioning a SQL Server Virtual Database (VDB)

1. Login to the **Delphix Management** application.
2. Select **Manage > Datasets**.
3. Select a **dSource**.
4. Click **Timeflow** tab.
5. Next to a snapshot select the  Provision VDB icon.
The **Provision VDB** panel will open, and the **Database Name** and **Recovery Model** will auto-populate with information from the dSource.
6. Select a **target environment**.
7. Select an **Instance** to use.
8. If the selected target environment is a Windows Failover Cluster environment, select a drive letter from **Available Drives**. This drive will contain volume mount points to Delphix storage.

Windows Cluster Volume Management Software Requirements

Only cluster volumes managed by the native Windows Volume Manager are supported. For example, cluster volumes managed by Veritas VxVM are not supported.

If you use third-party volume management software, create a new LU (recommended to be 10GB in size) and add this LU as a clustered resource to the SQL Server instance using native Windows volume management tools. Assign a drive letter for this LU. You can then use this LU as the volume mount point location for Delphix VDB provisioning.

9. Enter a **VDB Name** and select a **Target Group** for the VDB.
10. Enable **Auto VDB Restart** to allow the Delphix Engine to automatically restart the VDB when it detects target host reboot.
11. Click **Next**.
12. Select a **Snapshot Policy** for the VDB. Click **Next**.
13. Specify any **Pre-** or **Post-Scripts** that should be used during the provisioning process.
14. Click **Next**.

15. The final summary tab will enable you to review your configurations.
16. Click **Submit**.

When provisioning starts, the VDB will appear in the **Datasets** panel. Select the VDB and navigate to the **Status** tab to see the progress of the job. When provisioning is complete, you can see more information on the **Configuration** tab.

- ✔ You can select a SQL Server instance that has a higher version than the source database and the VDB will be automatically upgraded. For more information about compatibility between different versions of SQL Server, see [SQL Server Support Matrix](#).

Provisioning by Snapshot or LogSync

When provisioning by snapshot, you can provision to the start of any particular snapshot, either by time or LSN.

- ✔ You can take a new snapshot of the dSource and provision from it by clicking the **Camera** icon.

Provisioning By Snapshot	Description
Provision by Time	You can provision to the start of any snapshot by selecting that snapshot card from the TimeFlow tab, or by selecting  and entering a value in the time entry fields. The values you enter will snap to the beginning of the nearest snapshot.
Provision by LSN	You can use Provision by LSN control to open the LSN entry field. Here, you can type or paste in the LSN to which you want to provision. After entering a value, it will "snap" to the start of the closest appropriate snapshot. Provisioning a SQL Server VDB Procedure

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.
2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision new VDBs from that masked VDB to quickly provide safe data to the development and QA teams.

Related Topics

- [Linking Data Sources and Syncing Data with SQL Server](#)
- [Provisioning and Managing Virtual Databases with SQL Server](#)

SQL Server Support and Requirements

This section covers the following topics:

- [SQL Server Support Matrix](#)
- [Overview of Requirements for SQL Server Environments](#)
- [Requirements for Windows iSCSI Configuration](#)
- [Network Access Requirements for SQL Server](#)

SQL Server Support Matrix

 • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

• Delphix only supports 64-bit operating systems.

Key:

Color	Supported?
	Yes
	Not Applicable

Windows Server

Supported OS Version	Supported DBMS Version		
	SQL Server 2012	SQL Server 2014	SQL Server 2016
Windows Server 2012			
Windows Server 2012 R2			
Windows Server 2016			
Windows Server 2019			
Windows Server 2022			

 **Windows Support**

Delphix only supports "Windows with Desktop Experience" server Installation. Windows Core (Windows without Desktop Experience) is NOT included in our general support.

 **Windows Support**

After applying Windows updates or .NET updates for Windows Server 2016 or 2019, servers used to host Delphix VDBs may experience the following symptoms:

- CPU usage on the Windows target host may increase significantly, possibly reaching 100%.

- PowerShell may take up to 10 seconds to launch from the command line.

For more information on how to resolve this issue, see [KBA6024](#).

64-bit Windows only

Delphix supports only 64-bit versions of Windows on VDB Target Hosts and Staging Target Hosts. This restriction does not apply to Source Hosts.

Check OS Compatibility

The Windows Server OS versions on the Staging Target Hosts and VDB Target Hosts must be compatible. Please refer to the compatibility matrices below.

SQL Server Failover Clusters

There are additional restrictions on supported Windows and SQL Server versions for SQL Server Failover Cluster target environments.

For details, see [Adding a SQL Server Failover Cluster Target Environment](#).

Supported SQL Versions

SQL Server Version	Delphix Version
SQL Server 2012 (11.0)	Delphix 3.1.2 and beyond
SQL Server 2014 (12.0)	Delphix 4.1.3 and beyond
SQL Server 2016 (13.0)	Delphix 5.1.4 and beyond
SQL Server 2017 (14.0)	Delphix 5.2.3 and beyond
SQL Server 2019 (15.0)	Delphix 5.3.7 and beyond

Supported Windows Connector Versions

You can always download the current Windows connector version from the Delphix Engine from the following locations:

Delphix Dynamic Data Platform versions before 6.0.5.0:

```
http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.msi
```

Delphix Dynamic Data Platform versions 6.0.5.0 onwards:

```
http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.exe
```

Note

The Delphix Connector and Engine are generally backward-compatible. Delphix connector version 1.18.0 and beyond is compatible with Delphix Engine 6.0.5.0 and later releases. Delphix connector version 1.17.0 or earlier is compatible with Delphix Engine 6.0.4.0 and earlier releases.

Checking the Windows Connector Version

After installing the connector, you can perform the following steps to determine the Windows connector version.

1. Navigate to the Windows Environment.
2. In the search box on the taskbar, type Add or Remove Programs or launch **Start > Control Panel > Programs > Programs and Features**.
3. On the Add or Remove Programs screen, locate Delphix Connector and then click on the software to check the version installed.

Windows Connector Matrix

The following table details the supported Windows Connector and JRE versions for the corresponding Delphix Dynamic Data Platform versions.

Delphix Dynamic Data Platform Version	Windows Connector Version	JRE Version
4.2	1.2.0	1.7.0_71 (4.2.5.1)
4.3	1.3.0	1.8.0_40 (4.3.5.1)
5.1.0.0-5.1.4.0	1.4.0	1.8.0_40
5.1.5.0-5.1.9.0	1.5.0	1.8.0_112
5.2.0.0-5.2.2.0	1.9.0	1.8.0_131 (5.2.2.1)
5.2.3.0-5.2.4.0	1.10.0	1.8.0_162
5.2.5.0-5.3.3.0	1.12.0	1.8.0_202
5.3.4.0	1.13.0	1.8.0_202
5.3.5.0	1.14.0	1.8.0_202
5.3.6.0-5.3.7.0	1.15.0	1.8.0_202
5.3.8.0-6.0.0.0	1.16.0	1.8.0_202

Delphix Dynamic Data Platform Version	Windows Connector Version	JRE Version
6.0.1.0-6.0.4.0	1.17.0	OpenJDK 8u242
6.0.5.0-6.0.8.0	1.18.0	OpenJDK 8u262-b10
6.0.9.0-6.0.10.0	1.20.0	OpenJDK 8u282-b08
6.0.11.0	1.21.0.0	OpenJDK 8u302-b08
6.0.12.0	1.22.0.0	OpenJDK 8u302-b08
6.0.12.1	1.22.0.0	OpenJDK 8u302-b08
6.0.13.0	1.23.0.0	OpenJDK 8u302-b08

The following table details the .Net connector version compatibility for the corresponding Windows Connector versions.

.Net Connector Version	Windows Connector Version
3.5	1.17.0 and earlier
3.5 and 4.0+	1.18.0 and higher

File Checksum / Hashes

The following table provides the MD5 and SHA256 file hash for recent Delphix versions, in the instance that a security policy requires an exception for software installation.

Delphix Dynamic Data Platform Version	Connector Version	MD5 Hash	SHA256 Hash
6.0.0.0	1.16.0	84616c4f86c0d871127feb26b9b8bd52	bc2e5e4b26c33d739f7be12f53637979481e1371325e0b5c7645462f6cf5f4f3
6.0.1.0	1.17.0	bb6544dcff3f22d6bf9e20429b1cdb0c	8d914b77d2c5aa0695c3fdb7d7394e8ae051c091f6f5ee59185cb1d9d157542

Delphix Dynamic Data Platform Version	Connector Version	MD5 Hash	SHA256 Hash
6.0.1.1	1.17.0	dd3ea3bedf92a1f8d7e89228315ee315	75ed8aa2cf2689916e0b92de6c2aa2fa7660003d62ff07e99978f95dcc9f6abc
6.0.2.0	1.17.0	4770f8c1ad46430a1e4b8764ad687ff5	a98f1c1f90d5053f62d74c1a74bc27b3b48f68a3e067dce04ce9ab2ccf484241
6.0.2.1	1.17.0	97ec3d13db50ad843625b7757c1a2360	7c2ab497ae4ac27b907611e976d7068ed8eaf54e52b836591c082bbe01079512
6.0.3.0	1.17.0	3154c7a5fd7bcb4f0e2c66d328290e6a	efb69b2df884705e488d3405a3958c3bb5457142cbfc71990559bb549c2e8c37
6.0.3.1	1.17.0	2f73be2cf8ceef97cf519c716237462d	7fa7bc269bbc944568231f93ac17ba919b494aa43009967245b75f4abe576a73
6.0.4.0	1.17.0	6a4b7f0a5d180438ff199b2b464506cd	971b88456574e8e11906b5e6850ea46f012490dcb789a5af73f37d71158ffebf
6.0.4.1	1.17.0	ea9a564d0cae1066cd9607824292fbd7	674f9896360a351b314d6aa6e624f9270ef9e3e8d113baddf028581f1ddddec9
6.0.4.2	1.17.0	c5f855c1fa26f5cd3d84d70a94fc0480	07d71acafd559e15782e7145225734adf5e50356d2f8abe5a52bbd90f3a44cbe
6.0.5.0	1.18.0	a3ac768989a6c1dae41adf9452548236	71525c31edcf321d2c0220f4edf8932e249e2cd14b261054af1859574dc9e5eb
6.0.6.0	1.18.0	acb09cbd754e9c25d1afa17b35321cd4	3a51bd3b05148dc3563c5b36d83d44ac725295c23c4b7e9c28f54263058fe2b4
6.0.6.1	1.18.0	32202760adb743a62851f8b698a31bf3	ac16e8a5c37807c2fa775ecc18f9e0387f5bd9d235c9e8854275c99d5bf865a1
6.0.7.0	1.18.0	2fed8da4b986b95a7efb9a9a148239ba	2c2a83bad568617bf5dd36b84ad0de22b811f113acac45124ed0effbeb6aef17
6.0.8.0	1.18.0	5bf0b950cdcad8c39db60e1a7406d985	eae5958878a22b4f4f0456210c031ef9c11d064a05ca6cb85535f0e7dec9c0a1

Delphix Dynamic Data Platform Version	Connector Version	MD5 Hash	SHA256 Hash
6.0.8.1	1.18.0	a3ac768989a6c1dae41adf9452548236	71525c31edcf321d2c0220f4edf8932e249e2cd14b261054af1859574dc9e5eb
6.0.9.0	1.20.0	127af1d95372f84feb373ab4f7b280b2	b7c7aaceb2b04aebfbd0ba2521dd15520159d54877b0543567b892ee2584c7ce
6.0.10.0	1.20.0	e8a39a57e5a77d106d8040a4537bd8f0	7d4501bf5f1cf22a7039e8f7429d4b42c0d11b1360434380b174830f1155646e
6.0.10.1	1.20.0	baf1223b73e626798dd404bd4d18a3c9	b827d9213e933fce8389322fe25c8d2e1db57ab66709d236b9a8cd5e254bb89d
6.0.11.0	1.21.0.0	412c4c32c7aefb973c21b573dca66d2a	9b689a21249456ddf623042cd2595820531cfbfca0703e713d9795036eef49b7
6.0.12.0	1.22.0.0	dc7a97c46ef9e2d91285dcb0712eaf61	db544774ac3eb3235bf90deef3e93962f3955e86bd9f14309ecb0fba9d41dd7f
6.0.12.1	1.22.0.0	8cf9e955438c6393031c5596ca0420cd	130249658f68eab2a2263deceed6a9de6dd54eead66da68886272f3f40b2f9af
6.0.13.0	1.23.0.0	31cc5d3869e4c7f93d38014e5bd40b27	de1c2fd4c79dcb7391a51774fba1821456ecc031283db5b497145cce383005f2

SQL Server Editions

For Staging Target Hosts, you can use SQL Server Standard Edition only if the source databases do not use SQL Server Enterprise Edition specific features, such as partitioned tables.

However, it is important for VDB target hosts to use the same edition of SQL Server software as the source database so that all features available in the source are also available in the VDB.

Source Environment Operating System Compatibility

Source environments can be running any supported Windows operating system version. There are no compatibility requirements between the source environment's operating system and that of the target environments. For more information, see [Overview of Setting Up SQL Server Environments](#).

Staging and Target Environment Operating System Version Compatibility

The operating system version on the target environment that will contain the virtual databases must be equal to or higher than that on the staging target.

Source and Staging Environment SQL Server Version Compatibility

The SQL Server version on the staging environment must be equal to that on the source environment.

Provisioning to Higher SQL Versions When the Source is SQL Server 2005

For SQL Server 2005, direct provisioning to higher SQL Server versions is only supported for provisioning to SQL Server 2012 or higher. You can first provision a VDB to SQL Server 2005 and then upgrade it to a higher version by following the steps outlined in [Upgrading SQL Server VDBs](#).

Supported SQL Server Backup Software

Delphix currently supports the following software for dSource backups:

- SQL Server native backups - all formats, including but not limited to
 - Striped
 - Compressed
 - Append
 - Overwrite
- Veritas NetBackup
 - Delphix currently supports NetBackup v7.7.3, v8.0, and v8.1.
 - The version of the NetBackup client on the staging environment must be the same as that on the source.
 - If the dSource is backed up with NetBackup, the source and staging environments must have a NetBackup client installed.
 - VDB target environments do not need to have NetBackup installed.
 - Logsync (point-in-time provision) is currently not supported.
 - See [Linking a dSource from a NetBackup SQL Server Backup](#) for more NetBackup requirements and additional information about NetBackup.
- Commvault
 - Delphix currently supports Commvault v11.
 - The version of Commvault SQL Agent on the staging environment must be the same as that on the source.
 - If the dSource is backed up with Commvault, the source and staging environment must have Commvault SQL Agent installed.
 - VDB target environments do not need to have Commvault SQL Agent installed.
 - Commvault SQL Agent on staging environment must be registered to the CommServe server, which is managing source database backups.
 - Logsync (point-in-time provision) is currently not supported.
 - See [Linking a dSource from a Commvault SQL Server Backup](#) for more Commvault requirements and additional information about Commvault.
- Quest/NetVault LiteSpeed
 - Delphix currently supports LiteSpeed v5.0.0.0 to v8.9.
 - The LiteSpeed version on the staging environment must be the same or higher than that on the source.
 - If the dSource is backed up with LiteSpeed, the source and the staging environments must also have LiteSpeed installed.
 - VDB target environments do not need to have LiteSpeed installed.
 - On source SQL Server Instance, the dSource user must be granted execute privilege on extended stored procedure master.dbo.xp_sqlightspeed_version.
- Red Gate SQL Backup Pro
 - Delphix currently supports SQL Backup Pro v7.3 to v7.6 and v10.0.
 - The version of SQL Backup Pro on the staging environment must be the same as that on the source.

- On source SQL Server Instance, the dSource database user must be granted execute privilege on extended stored procedure master.dbo.sqbutility.

If the dSource is backed up with SQL Backup Pro, the source and the staging environments must have SQL Backup Pro installed.

VDB target environments do not need to have SQL Backup Pro.

Source and Provisioning Environment SQL Server Compatibility Matrix

When provisioning a VDB, the SQL Server version on the target should be equal to or higher than that on the source.

	Provisioning Target Environment				
	SQL Server 2012	SQL Server 2014	SQL Server 2016	SQL Server 2017	SQL Server 2019
Source Environment					
SQL Server 2012	X	X	X	X	X
SQL Server 2014		X	X	X	X
SQL Server 2016			X	X	X
SQL Server 2017				X	X
SQL Server 2019					X

Related Topics

- [Linking a dSource from a Commvault SQL Server Backup](#)
- [Linking a dSource from a NetBackup SQL Server Backup](#)
- [Upgrading SQL Server VDBs](#)


Overview of Requirements for SQL Server Environments

SQL Server Staging Hosts and Databases

This document identifies the requirements for interactions between the Delphix Engine and SQL Server environments and outlines the set of system tables to which we currently require access.


Delphix SQL Server Architectural Diagram - Traditional Pull Architecture

This diagram depicts the environments and hosts with which the Delphix Engine interacts. Each type of environment has different requirements, which are described below.



Delphix SQL Server Architectural Diagram - Staging Push Architecture

This diagram depicts the Staging and target hosts with which the Delphix Engine interacts. Note that the Delphix Engine does not interact with the Source Database. For more details on the Staging Push mechanism, see [Staging Push Mechanism with SQL Server](#) and [Staging Push Implementation for SQL Server](#).



SQL Server Source Hosts and Databases

Source Host Requirements

Windows servers that will be added as Source Environments must meet the following requirements:

Source Host Requirement	Explanation
<p>To allow Delphix-initiated backups, the service account running each SQL Server instance (the Instance Owner) should be one of:</p> <ul style="list-style-type: none"> • A domain user (e.g. <code>MYDOMAIN\accountname</code>) (RECOMMENDED) • A Managed Service Account or Group Managed Service Account (<code>MYDOMAIN\accountname\$</code>) (requires Windows 2012 and later, and SQL Server 2008R2 or later) • The LOCAL SYSTEM account (<code>NT AUTHORITY\SYSTEM</code>) • The NETWORK SERVICE account (<code>NT AUTHORITY\NETWORK SERVICE</code>) 	<p>Backups initiated by the Delphix Engine (Delphix Managed Backups or manual snapshots which request a backup) will fail if the Delphix Engine cannot map the Instance Owner to an Active Directory user or computer object.</p>
<p>The source host, proxy, and staging environments must have appropriate cross-domain trust relationships</p>	<p>For more information on these requirements, see the document Delphix in Multi-domain Windows Environments.</p>

Source Windows User Requirements

Authenticate to Active Directory

Authenticate to Active Directory using a username and password: Delphix authenticates to Active Directory using a username and password. The Windows domain user cannot authenticate using Managed Service Accounts, Group Managed Service Accounts, smart cards, or other authentication methods. The Delphix Engine can retrieve passwords from a supported password vault, as described in [Password Vault Support](#).

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding source environments to the Delphix Engine. The Windows domain user provides remote read-only access to the Windows Registry, which allows the discovery of SQL Server instances and databases (note: Delphix also uses the registry to discover cluster nodes in AG Cluster Source Environments). This user must have the following permissions:

Source Windows User Requirement	Explanation
<p>Be a member of the Backup Operators group on the Source Host</p>	<p>Delphix uses this privilege to discover SQL Server instances and gather system details, using Windows remote registry access.</p>

Source Windows User Requirement	Explanation
Be a member of the Users group on the Staging Host	In order to discover and query SQL Server instances as the Source Windows User, scripts are run on that user at the Staging Host.
Have the "Log on as batch" privilege on the Staging host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Source Database Login Requirements</i> section below.

Source Database Login Requirements

The Delphix Engine requires SQL Server logins to be created on each SQL Server instance that the Delphix Engine will communicate with:

- A database login for Environment discovery and monitoring, specified when Adding an Environment. This must be a Windows Authentication login for the Source Windows User configured in the previous section.
- A database login for dSource (Source Database) monitoring and interaction, specified when Linking a dSource. This user can be:
 - The same as the Source Windows User;
 - A different Windows Authentication login (this user must also have Log on as a Batch Job privileges on the Staging host); or
 - A SQL Authentication login

These users must have the following permissions on each instance:

Object	Privileges Required	dSource User	Environment User	Purpose
Server	CONNECT SQL	✓	✓	Access to the SQL Server instance
Database: master	db_datareader	✓	✓	Access to information about attached databases
Database: msdb	db_datareader	✓		Access to the backup history
Each user database to be linked	PUBLIC	✓		Delphix will periodically run queries to check the current size of the database

Object	Privileges Required	dSource User	Environment User	Purpose
Each user database to be linked	db_backupoperator	✓		Optional: Required for backups to be initiated by Delphix (using Delphix Managed Backups, or when Delphix initiates a backup during a manual Snapshot)
Server	VIEW ANY DEFINITION	✓	✓	Optional: Required for the discovery of databases in Availability Groups
Server	VIEW SERVER STATE	✓	✓	Optional: Required for the SnapSync and discovery of Availability Groups
Object: master.dbo.sqlbutility	EXECUTE	✓		Optional: Required when using backups created by Red Gate SQL Backup
Object: master.dbo.xp_sqlightspeed_version	EXECUTE	✓		Optional: Required when using backups created by Quest LiteSpeed for SQL Server

List of Source tables accessed by the Delphix Engine

Using the db_datareader permission, the Delphix Engine accesses the following system tables in the master and msdb databases on the source host:

System table	Justification
master.sys.databases	Used to determine the name and recovery model of databases within discover SQL Server instances
master.sys.availability_groups	Used for discovering all the availability groups within an Availability Group source environment.

System table	Justification
master.sys.availability_group_listeners	Used for discovering all the availability group listeners within an Availability Group source environment. <ul style="list-style-type: none"> A requirement for dSource linking of SQL Server clustered databases (replicas) is to provide an AG (Availability Group) listener for AG cluster source discovery. This is implemented on AG cluster source discovery as a failsafe if AG cluster source database authentication configuration change down the line, ensuring the Delphix engine has a way to reach the cluster and continue certain operations.
master.sys.availability_databases_cluster	Used for discovering all the availability group clusters within an Availability Group source environment.
master.sys.availability_replicas	Used for discovering all the availability group replicas within an Availability Group source environment.
master.sys.database_files	Used to determine the size of databases and whether filestream is enabled for a database
master.sys.dm_exec_requests	Used to enable Delphix to report backup operation progress
master.sys.master_files	Used to determine the primary file of a database
master.sys.filegroups	Used to determine the filegroups of a database so that Delphix can configure VDBs with the same filegroups
msdb.dbo.backupset	Used to determine new backups that have been taken. This table is regularly queried to find out if a new backup image has been taken and needs to be synchronized with Delphix.
msdb.dbo.backupmediafamily	Used to determine the physical device names of the backup files comprising a backup.

SQL Server Staging Hosts and Databases

Staging Host Requirements

Windows servers which will be added as Staging Environments must meet the following requirements:

Staging Host Requirement	Explanation
<p>The service account running each SQL Server instance (the Instance Owner) must be one of:</p> <ul style="list-style-type: none"> • A domain user (e.g. <code>MYDOMAIN\accountname</code>) (RECOMMENDED) • A Managed Service Account or Group Managed Service Account (<code>MYDOMAIN\accountname\$</code>) (requires Windows 2012 and later, and SQL Server 2008R2 or later) • The LOCAL SYSTEM account (<code>NT AUTHORITY\SYSTEM</code>) • The NETWORK SERVICE account (<code>NT AUTHORITY\NETWORK SERVICE</code>) 	<p>Access to existing database backups for Snapshots and Validated Sync operations will fail if the service account cannot access backups created by the Source database instance.</p>
<p>The source host, proxy and staging environments must have appropriate cross-domain trust relationships</p>	<p>For more information on these requirements, see the document Delphix in Multi-domain Windows Environments.</p>
<p>The edition of the installed SQL Server instance(s) must support all database features used by linked Source databases</p>	<p>SQL Server may raise an error during some dSource operations if the Staging SQL Server Instance does not support features used by the Source Database.</p> <p>This is most easily addressed by using the same edition of SQL Server as the Source database.</p> <p>Features in use on the Source database can be checked using the <code>sys.dm_db_persisted_sku_features</code> dynamic view.</p>
<p>Must have both the Source and Staging Windows Users configured as local Windows users</p>	<p>See the sections <i>Source Windows User Requirements</i> and <i>Staging Windows User Requirements</i> for more detail.</p>
<p>Delphix Connector software is installed and running</p>	<p>See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.</p>
<p>Recommended iSCSI Registry settings must be in place</p>	<p>See Requirements for Windows iSCSI Configuration and Knowledge Base Article KBA1251 for instructions on applying these settings*.</p> <p>*Please note that in 6.0.x, these performance settings are subject to change based on further tuning efforts by Delphix Engineering.</p>

Staging Host Requirement	Explanation
sqlcmd command line utility must be installed on the servers hosting VDBs and staging databases. The utility should be in the Delphix operating system user's PATH environment variable.	Delphix does not run sqlcmd (or any other process) directly on the source SQL Server instance. It runs sqlcmd on the staging/connector host.

Staging Windows User Requirements

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding Staging environments to the Delphix Engine. This user must have the following permissions:

Staging User Requirement	Explanation
Be a member of the Windows "Local Administrators" group on the Staging Host	We require this permission for mounting iSCSI LUNs presented by the Delphix Engine to the staging and target hosts. Microsoft utilities used by the Delphix Engine, such as diskpart, require membership of this group .
Have the "Log on as batch" privilege on the Staging host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to access existing backups from the Source Database	If accessing existing native, Red Gate, or LiteSpeed backups, the following permissions are required: <ul style="list-style-type: none"> The Staging Windows User and the service account running SQL Server must both have permission to access the database backups via SMB (Windows file sharing) The Staging Windows User and the service account running SQL Server must both have NTFS Permissions to access the database backups Separate documents describe requirements if Linking a dSource from a NetBackup SQL Server Backup or Linking a dSource from a Commvault SQL Server Backup .
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Staging Database Login Requirements</i> section below.

Staging Database Login Requirements

The Delphix Engine requires a Windows Authentication login to be created for the Staging Windows User on each SQL Server instance that the Delphix Engine will communicate with, with the following permissions:

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	CONNECT SQL	✓	Access to the SQL Server instance
Server	sysadmin	✓	<p>The staging and target databases are managed and administered completely by Delphix. Our functionality requires many administrative operations on those databases and requires full access to them.</p> <p>Since database ownership can be changed by customers as part of configuring virtual databases, we require the sysadmin role to continue to administer the databases.</p>

SQL Server Target Hosts and Databases

Target Host Requirements

Windows servers which will be added as Target Environments must meet the following requirements:

Target Host Requirement	Explanation
Delphix Connector software is installed and running	See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.
Recommended iSCSI Registry settings must be in place	See Requirements for Windows iSCSI Configuration and Knowledge Base Article KBA1251 for instructions on applying these settings.
The edition of the installed SQL Server instance(s) must support all database features used by linked Source databases	<p>SQL Server may raise an error during some dSource operations if the Staging SQL Server Instance does not support features used by the Source Database.</p> <p>This is most easily addressed by using the same edition of SQL Server as the Source database.</p> <p>Features in use on the Source database can be checked using the <code>sys.dm_db_persisted_sku_features</code> dynamic view.</p>

Target Host Requirement	Explanation
sqlcmd command line utility must be installed on the servers hosting VDBs and staging databases. The utility should be in the Delphix operating system user's PATH environment variable.	Delphix does not run sqlcmd (or any other process) directly on the source SQL Server instance. It runs sqlcmd on the staging/connector host.

Target Windows User Requirements

The Delphix Engine needs a Windows domain user — for example, `MYDOMAIN\delphix_os` — to be specified when adding Target environments to the Delphix Engine. This user must have the following permissions:

Target User Requirement	Explanation
Be a member of the Windows "Local Administrators" group on the Target Host	We require this permission for mounting iSCSI LUNs presented by the Delphix Engine to the staging and target hosts. Microsoft utilities used by the Delphix Engine, such as diskpart, require membership of this group .
Have the "Log on as batch" privilege on the Target Host	We require this permission for remote PowerShell execution. This privilege can be assigned through the Local Security Policy (Local Policies → User Rights Assignment → Log on as batch job).
Be able to login to each SQL Server instance that the Delphix Engine will communicate with	These requirements are described in the <i>Target Database Login Requirements</i> section below.
Delphix Connector software is installed and running	See Installing the Delphix Connector Service on the Target Database Servers for instructions on installing the Delphix Connector.

Target Database Login Requirements

The Delphix Engine requires a Windows Authentication login to be created for the Target Windows User on each SQL Server instance that the Delphix Engine will communicate with, with the following permissions:

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	CONNECT SQL		Access to the SQL Server instance

Object	Privileges Required	Delphix OS User (Windows Login)	Purpose
Server	sysadmin	✓	<p>The staging and target databases are managed and administered completely by Delphix. Our functionality requires many administrative operations on those databases and requires full access to them.</p> <p>Since database ownership can be changed by customers as part of configuring virtual databases, we require the sysadmin role to continue to administer the databases.</p>

Supported Roles for Failover Cluster Instances and Always On Availability Groups

Failover Cluster Instances and Always On Availability groups cannot be used as Staging Environments, and VDBs cannot be provisioned into Availability Groups.

 When adding a Failover Cluster Instance or Always On Availability Group, all nodes of the cluster must meet the requirements described in this document.

The following table shows how different SQL Server instance types should be added from the Add Environment screen:

Color		Supported?		
✓		Yes		
✗		No		
		Environment Role		
Instance Type	Added As	Source Environment	Staging Environment	Target Environment
Standalone Instance	Standalone	✓ *	✓	✓
Failover Cluster Instance (FCI)	Standalone	✓ *	✗	✗
Failover Cluster Instance (FCI)	Cluster	✗	✗	✓

		Environment Role		
Always On Availability Group (AG)	Cluster	Supported	✘	✘**

* Databases that are participating in Availability Groups will not be discovered during the discovery of a Standalone environment.

** VDBs cannot be provisioned into availability groups. However, SQL Server instances that participate in an Availability Group can also be added as Standalone Instances.

Using a Failover Cluster Instance as both Source and Target

A Failover Cluster Instance added as an environment once (as either a Source or Target environment) cannot be used as both a Source and Target.

If this is required, the environment can be added twice:

- Once as a Standalone Source environment
- Once as a Cluster Target environment

The Standalone Source environment can be used for linking dSources, and the Cluster Target environment is used for provisioning VDBs.

As suggested by the Best Practice note earlier in this article, this is not a recommended configuration.

Where possible, SQL Server failover cluster instances that the Delphix Engine will use as a target should not be used to host databases other than Delphix VDBs.

Requirements for Failover Cluster Target Environments

The following additional requirements exist for Windows Failover Cluster Instances added as Target Environments, to be used for VDB Provisioning.

- You must first add each node in the Window Failover Cluster individually as a standalone target environment, using a non-clustered address. See [Adding a SQL Server Standalone Target Environment](#).
 - The Delphix Engine may show a warning that it cannot discover the Failover Cluster Instances on each standalone host. This is expected.
- Each clustered SQL Server instance must have at least one clustered disk added to the clustered instance resource group, which can be used for creating mount points to Delphix storage.
 - The clustered drive must have a drive letter assigned to it.
 - The clustered drive must be formatted using the "GUID Partition Table (GPT)" partition style, in order for the Delphix Engine to automatically discover the drive letter as a valid option for the cluster instance. An MBR-formatted disk requires manual verification outside of Delphix that the disk has been correctly added to the MSSQL clustered resource group prior to creating the VDB. When provisioning the VDB, you must manually specify the desired MBR disk, because it will not appear in the Delphix GUI.
 - The clustered drive must be added to the clustered instance resource group as a dependency in the Failover Cluster Manager.
- Each node in the cluster must have the Failover Cluster Module for Windows PowerShell feature installed.
 - While running Windows PowerShell as an administrator, enter this command to test that the module is available: `Import-Module FailoverClusters`
- An additional target environment that can be used as a Connector Environment must exist.
 - This environment must **not** be a node in the cluster.
 - This environment should be part of the same Active Directory domain as the cluster.

Additional requirements for Azure SQL Server Availability Groups

Microsoft's [tutorial](#) and deployment template for building a *SQL Server AlwaysOn Cluster in Azure* does not include the full range of network connectivity that is available in on-premise deployments.

In addition to the connectivity provided by this template:

- The Staging Server must be able to connect to the **Cluster IP Address** using the RPC / Remote Registry port **TCP 445**; and
- The Availability Group Listener must be configured with a TCP Port, such as the SQL Server default 1433


Connectivity to the Cluster IP Address can be tested using the following PowerShell command. This command should be run from the Staging Server, and specify the Cluster Hostname:

```
New-Object System.Net.Sockets.TcpClient("aodns-fc.mydomain.local", "445")
```

⚠ For Azure clusters, Delphix does not support DNN (Distributed Network Name) at the WSFC level. However, DNN (Distributed Network Name) is supported for creating a SQL AG listener in SQL clustering.

To allow this connectivity on the marketplace template, the following additional steps may be required:

1. On a node of the SQL Server cluster, use PowerShell to define a "probe port". The active cluster node will open this port so that the Azure Load Balancer can detect it.
Get-ClusterResource "Cluster IP Address" | Set-ClusterParameter -Name "ProbePort" -Value 59998
2. From the Failover Cluster Manager, ensure that the Core Cluster Resource has a valid IP address for its subnet (this may be configured with an invalid IP address such as 169.254.1.1).
3. On all nodes of the SQL Server cluster, ensure that the Windows Firewall allows inbound TCP traffic on TCP port 59998.
4. Extend the Azure Load Balancer (sqlLoadBalancer) configuration to route connections to the Cluster IP Address.
 - a. Add a Frontend IP Address for the Cluster IP Address, with an IP address matching that configured for the cluster.
 - b. Add a Health Probe for TCP port 59998.
 - c. Add a Load Balancing Rule for Port 445, using the newly created Frontend IP Address and Health Probe. More complicated Azure AlwaysOn deployments, such as those including multiple networks, may require additional steps to allow this connectivity.
5. The default Availability Group Listener is not configured with a TCP Port. Current versions of Delphix require this to be configured with a valid port number, such as the default port 1433. This can be changed via the AlwaysOn object tree in SQL Server Management Studio:



Related Topics

- [Adding a SQL Server Standalone Target Environment](#)
- [SQL Server Support and Requirements](#)
- [Delphix in Multi-domain Windows Environments](#)
- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Requirements for Windows iSCSI Configuration](#)
- [Registry Settings for Optimal Database Performance and Stability \(KBA1251\)](#)

Requirements for Windows iSCSI Configuration

Windows iSCSI configuration requirements are split into two types. These requirements are needed on both staging and target servers.

1. iSCSI configuration required for operational stability.
2. iSCSI configuration specific for performance improvement.

 When target environments are discovered, Delphix will configure the Microsoft iSCSI Initiator Service for Automatic startup.

iSCSI Configuration Required for Operational Stability

We require the following Microsoft iSCSI Initiator configuration parameters for Target and Staging Hosts. For details about configuring registry settings, see [How to Modify the Windows Registry](#) on the Microsoft Support site.


A Windows Server reboot is required after changing iSCSI configuration parameters.

Registry Key	Registry Value	Type	Data
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\iSCSI\Discovery	MaxRequestHoldTime	REG_DWORD	0x384 (900)
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk	TimeOutValue	REG_DWORD	0x384 (900)
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters\Interfaces\<Interface GUID>	TcpAckFrequency	REG_DWORD	0x1 (1)
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\Parameters	iSCSIDisableNagle	REG_DWORD	0x1 (1)
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}\<Instance Number>\<Parameters>	MaxRequestHoldTime	REG_DWORD	0x12C (300)

Identifying the Instance Number for iSCSI Control Class Initiator Drivers

1. From the Windows toolbar, click **Start** and select **Run** from the menu.
2. Type regedit in the **Open** field and click **OK**.
3. Go to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-

BFC1-08002BE10318}\<Instance Number> where the value of <Instance Number> is the one that shows a **DriverDesc** value of **Microsoft iSCSI Initiator**. Under the registry key, locate and expand the plus (+) sign next to instance number. In this example below the value of instance number is 0002.



1 Value of Instance Number

Related Topics

- [SQL Server Support and Requirements](#)

Windows iSCSI Configuration and Limits for Target and Staging Hosts

Windows Limitation

Windows supports up to 255 iSCSI LUNs maximum. This creates a hard limit on the number of VDBs that can be created because each VDB requires one or more iSCSI connections.

iSCSI connections - Staging

1. dSource linked with Logsync disabled = 1 LUN (DATA)
2. dSource linked with Logsync enabled = 2 LUNs (DATA and ARCHIVE)
3. dSource linked with Logsync disabled and SnapShot started (new COPY ONLY FULL BACKUP) = 3 LUNs (DATA and TEMP)
 - a. Once the SnapShot is completed the TEMP LUN will be destroyed and 1 LUN used
4. dSource linked with Logsync enabled and SnapShot started (new COPY ONLY FULL BACKUP) = 3 LUNs (DATA, ARCHIVE and TEMP)
 - a. Once the SnapShot is completed the TEMP LUN will be destroyed and 2 LUNs used
5. A maximum of ~120 dSources per Staging Target is recommended, assuming an average of 2 LUNs per source, which would mean 240 LUNs would be consumed for normal operation.
 - a. The proposed scenario would leave 13 additional iSCSI connections available for COPY ONLY FULL BACKUPS
 - b. For dedicated staging hosts, we do NOT use a Powershell process for monitoring.

iSCSI connections - Targets

1. VDB normal operations = 1 LUN (DATA)
 - a. No extra mounts required for SnapShot restore or refresh from source Snapsync
2. VDB point-in-time log actions (such as restore, refresh or provision from logs) = 2 LUNs (DATA and SOURCE_ARCHIVE)
 - a. An extra LUN is not required for Snapsync operations, only Logsync
 - b. Most users do not require enablement of the Logsync feature for MSSQL Sources, because sources in FULL RECOVERY mode create a Snapsync for each log file, providing a significant number of restore points even without retaining the logs.
 - c. Once recovery is completed the SOURCE_ARCHIVE LUN will be destroyed and 1 LUN used
3. A maximum of ~120 VDB's per Target is recommended
 - a. In 4.x, target host iSCSI connections are less likely to be a limitation while processing costs for Powershell threads may become prohibitive because each target VDB requires a Powershell process for monitoring
 - b. In 5.x, this has been alleviated with a hard limit on Powershell processes

iSCSI connections - V2P

1. V2P normal operation = 1 LUN (DATA)
 - a. Once the V2P operation is completed the DATA LUN will be destroyed leaving no LUNs used
2. V2P point in time log actions (provision from logs) = 2 LUNs (DATA and SOURCE_ARCHIVE)
 - a. Once the V2P operation is completed both the DATA and SOURCE_ARCHIVE LUNs will be destroyed leaving no LUNs used

Related Topics


- [Recommendations for Windows iSCSI Configuration](#)

Network Access Requirements for SQL Server

Network Architecture

The diagram *Delphix Virtualizing and Masking SQL Server Network Architecture* below depicts the overall network architecture for Delphix virtualizing and masking SQL Server. In the diagram, each of the arrows represents the direction of a network connection between two nodes. Next to each arrow is a label indicating the network protocol (TCP) and the port number indicating the network service. Also, indicated in red are the recommended network latencies between the major components of the architecture.

✔ You can optionally configure a separate Connector Environment, specifically used to discover databases on the source during Environment creation. You can also use your Staging Target Host to be used as the Connector Environment, as seen in the image below.



Delphix Virtualizing and Masking SQL Server Network Architecture

Ports

Based on the table below, the Windows Network Administrator needs to complete a series of tasks. For each port listed, determine whether it must be opened in your firewall between your Delphix Engine and source or target

systems. Work with your Delphix Administrator to understand what requirements are there, and ensure that they have been met before proceeding.

Port	Network Service	Required for virtualization?	Required for masking?	Description and usage
22	SSH	Yes	Yes	Used for accessing command-line interface (CLI) and internal Delphix OS accounts
80	HTTP	Yes	No	Used for GUI console access on Delphix Engine by default, disabled when HTTPS in use
443	HTTPS	Yes	No	Used for GUI console access on Delphix Engine, disabled when HTTP in use
445	SMB	Yes	No	Used for attaching shared folders on Windows. To take a copy-only backup or use Delphix Managed backups, this port is required to allow the source environment access to the staging environment.
1433	JDBC	Yes	Yes	Used for accessing SQL Server databases for queries on data-dictionary. This port is the default, but you can use other ports instead.
3260	iSCSI	Yes	No	Used for network-attached storage (NAS) on Windows database servers
8415	DSP	Yes	No	Used for SQL Server hooks, when enabled. This port needs to be open between the Delphix Engine and the VDB target.
(32768-60999)/9100	Delphix Windows Connector	Yes	No	Used for connecting to the Delphix Connector service installed on Windows target database servers.

Port	Network Service	Required for virtualization?	Required for masking?	Description and usage
50001	iPERF	No	No	Used for network throughput testing with the open-source iPerf package through the Delphix CLI, this is purely optional (but useful) functionality

Related Topics

- [SQL Server Support and Requirements](#)

Managing SQL Server Environments and Hosts


This section describes the attributes of SQL Server environments such as information for the Delphix Connector and covers the following topics:

- [Overview of Setting Up SQL Server Environments](#)
- [Environment Attributes for Hosts with SQL Server](#)
- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Manual Discovery for SQL Server Instances](#)
- [Adding a SQL Server Source Environment](#)
- [Adding a SQL Server Standalone Target Environment](#)
- [Adding a SQL Server Failover Cluster Target Environment](#)
- [Additional SQL Server Environment Topics](#)

Overview of Setting Up SQL Server Environments


This topic describes the high-level process for adding SQL Server environments, linking SQL Server databases to the Delphix Engine, and provisioning virtual databases.

Block Diagram of Linking Architecture between SQL Server Environments and the Delphix Engine



This block diagram shows the host architecture when creating SQL Server Environments in the Delphix Engine. On the left, the Source Host (in this example a Production Database Server) creates backups to a shared backup location. The Delphix Engine (far right), continuously monitors the source database to determine when new backups are available. When a backup is available, the Delphix Engine will contact the Windows Staging host (center) via the Delphix Connector service. The Staging host will read the backups from the shared backup location and recover them through a Staging Database that is automatically set up on Delphix iSCSI storage. Once recovery is complete, the backup data is incorporated into the Delphix dSource as a new snapshot card, and is available for use.

Block Diagram of SQL Server Provisioning Architecture



This block diagram shows how Delphix provisions VDBs to a Target Environment. The Delphix Engine (left) creates a set of virtual files from a snapshot that becomes the VDB container. These files are presented to the Target host (right) via iSCSI. Delphix calls the Delphix Connector on the Target and initiates the creation of a new database using the Virtual files. Once complete, the Virtual Database is brought online and made available for use.

The Delphix Connector and Environment Setup

The Delphix Connector is a Windows service that enables the communication between the Delphix Engine and the Windows target environment where it is installed.

- The directory on which you install the Delphix Connector should have at least 1GB of available space.

This target machine can serve three purposes in a Delphix Engine deployment. It can:

1. Serve as a proxy for database discovery on source hosts.
2. Host a staging database for a linked dSource and run the validated sync process.
3. Host a target environment for provisioning virtual databases (VDBs).

Database discovery is initiated during the environment set up process. When you specify a production source environment that contains the databases you want to manage with the Delphix Engine, you must also specify a target environment where you have installed the Delphix Connector to act as a proxy for communication with the source environment. This is necessary because Delphix does not require that you install the Delphix Connector software on the production source environment. When you register the source environment with the Delphix Engine, the Delphix Engine uses the Delphix Connector on the proxy environment to discover SQL Server instances and databases on the source. You can then create dSources from the discovered databases. If you later refresh the source environment, the Delphix Engine will execute instance and database re-discovery through the proxy host.

SQL Server dSources are backed by a staging database that runs on a target host. There is no requirement for additional local storage on this target host, as the storage is mounted over iSCSI from the Delphix Engine. At Delphix, we refer to the creation and maintenance of this staging database on the staging host as "validated sync," because it prepares the dSource data on the Delphix Engine for provisioning VDBs later on. After the Delphix Engine creates the staging database, it continuously monitors the source database for new transaction log backups. When it detects a new transaction log backup, it restores that backup to the staging database. The result is a TimeFlow with consistent points from which you can provision a VDB, and a faster provisioning process, because there is no need for any database recovery during provisioning.

When you later provision a VDB, you can specify any environment as a target, including the environment that contains the staging database. However, for best performance, we recommend that you choose a different target environment. The only requirements for the target are:

- It must have the Delphix Connector installed.
- It must have an operating system that is compatible with the one running on the validated host, as described in [SQL Server Support and Requirements](#).

Toolkit Size and Predicted Growth

Each of the clients that run from the client-side toolkit generates its own logs. Each client generates 4 different log files, one for each level of logging: info, trace, debug, and error. Each level of logging is restricted to a maximum of 10 logfiles, and these logfiles are capped at 10MB each. Therefore, the Delphix Engine will consume a maximum of 400MB per client-side application. On the source server, there are currently two commonly run client-side applications, SnapSyncClient and the Delphix Connector.

✔ V2P also generates its own logs, so if you intend to V2P to the source, you should account for an additional 400MB in your upper bound.

Thus, the max amount of growth for the toolkit from logging is 800MB without V2P (or 1.2GB with V2P).

Linking additional dSources does not impact the size of the toolkit on production, aside from the log messages generated during linking, which is included in the calculation above.

On the target server, unlike the source server, there would be only one client – the Delphix Connector, which would occupy around 400 MB maximum storage space. In addition, the Delphix Engine pushes new scripts each time you provision a VDB. This requires < 1MB space.

Therefore, the maximum space occupied by the toolkit directory on the source server is its initial size (~ 400MB) + 800MB = 1.2 GB. While on the target server, the maximum toolkit size is initial size (~ 400MB) + 400 MB + Number of VDBs * 1MB.

Related Topics

- [SQL Server Support and Requirements](#)
- [SQL Server Environments and Data Sources](#)

Environment Attributes for Hosts with SQL Server

This topic describes the attributes of SQL Server environments such as information for the Delphix Connector. Below you will see a section for common environment attributes shared by all types of environments as well as SQL Server-specific ones.

Procedure

1. Login to the Delphix Management application.
2. Click Manage.
3. Select Environments.
4. In the Environments panel, click the name of an environment to view its attributes.
5. Next to Attributes, click the Pencil icon to edit an attribute.

Common Environment Attributes

Attribute	Description
Environment Users	The users for that environment. These are the users who have permission to ssh into an environment or access the environment through the Delphix Connector. For more information on the environment user requirements, see the Requirements topics for specific data platforms.
Host Address	The IP address of the environment host.
DSP KeyStore Path	The path to the user-managed DSP Keystore.
DSP KeyStore Alias	The lowercase alias to use inside the user-managed DSP Keystore.
DSP KeyStore Password	The password for the user-managed DSP Keystore.
DSP TrustStore Path	The path to the user managed DSP truststore.
DSP TrustStore Password	The password for the user managed DSP truststore.
OS	The name of the host operating system.

Attribute	Description
Version	The version of the host operating system.
Release	The release of the host operating system.
Time Zone	The timezone of the host operating system.
Total RAM	The amount of RAM on the host machine.
Processor Type	The processor type of the host machine.
Traceroute	Traceroute info from target host to Delphix Engine.
Notes	Any other information you want to add about the environment.
Java Development Kit	The currently selected JDK kit will be shown.
Java Development Kit (JDK) Path	Location of the Java Development Kit (JDK) used for the host. Only specified if the feature to provide your own JDK is enabled, otherwise, the defaults are used per our Java Support Matrix .

SQL Server Environment Attributes

Attribute	Description
Installed Powershell Version	The PowerShell version installed on the windows target host. This is not applicable for Windows source environments.
Delphix Connector Path	The path for the toolkit that resides on the host.
Delphix Connector Port	The port that the connector connects on.
Delphix Connector Version	The Windows Connector version that is installed on the provided host.
.NET Framework Version	The .NET Framework version used for Windows Connector Service.

Delphix Connector Host	The connector host for a windows environment. This is not applicable for Windows target environments.
------------------------	---

Related Topics

- [Managing SQL Server Environments and Hosts](#)

Installing the Delphix Connector Service on the Target Database Servers

This section lists the steps involved in installing a Delphix Connector on your target database server. Installing the Delphix Connector is vital for communication between the Delphix Engine and the targets. A minimum available space of 1GB is a prerequisite to installing the Delphix Connector.

1. From the machine that you want to use as a target, start a browser session and connect to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Plus** icon.
5. In the **Add Environment** dialog, select **Windows** from the **operating system** menu.
6. Select **Target**.
7. Select **Standalone**.
8. Click **Next**.
9. Click the download link for the **Delphix Connector Installer**.
The Delphix Connector will download to your local machine.

 You can also download the Delphix Connector directly by navigating to this URL: `http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.exe`

10. On the Windows machine that you want to use as a target, run the Delphix Connector installer. Click **Next** to advance through each of the installation wizard screens.

 The installer will only run on 64-bit Windows systems. 32-bit systems are not supported.

- a. For **Connector Configuration**, make sure there is no firewall in your environment blocking traffic to the port on the target environment that the Delphix Connector service will listen to.
- b. For **Java Configuration**, to provide your own Oracle Java enter the absolute path to your Oracle JDK and click **Next**. Otherwise, leave the field blank.

 The NETWORK SERVICE user requires read and execute permissions on the Oracle JDK, its subfolders, and files.

- c. For **Select Installation Folder**, either accept the default folder or click **Browse** to select another.
- d. For **Choose .NET Framework (Delphix Engine 6.0.5.0 onwards)**, select which version of .NET should be used by the Delphix Connector Windows service. The installer will default to the latest version detected on the system. This selection can be changed afterwards but make sure that chosen .NET framework version should exist in the system.
- e. Click **Next** on the installer final **Confirm Installation** dialog to complete the installation process.
- f. For successful installation a popup will come up stating "**DelphixConnector Installed Successfully**".
- g. Click **Close** to exit the Delphix Connector Install Program.

Related Topics

- [Relocating the Delphix Connector](#)
- [Replacing Self-signed Certificates on the Delphix Connector](#)

Relocating the Delphix Connector

There are times when the Delphix Connector installation requires a move to a different directory or drive. It's not a trivial relocation. This requires disabling **dSources** and/or **VDBs**, uninstalling the current install and reinstalling to the new location. In addition to this, an upgrade to the Delphix Connector can also be achieved via the uninstall/reinstall methodology, including a change in location. These instructions cover 4.0 through 5.1. The steps are the same up to 5.1.2.0. At that version and forward it is no longer required to use the CLI to change the Delphix Connector location. All one needs to do is refresh the Delphix Connector environment and the new directory location is discovered and updated on the Delphix Engine. In the steps listed below, the example is moving the connector from "C:\Program Files" to "C:\", so the full connector path is "C:\Delphix\DelphixConnector".

1. **Disable the **dSources** and/or **VDBs** associated to the **DelphixConnector** host.**
This will unmount the storage from the Windows host, removing the directories, the **dSources** represented by the staging databases, and the **VDBs**.
 - a. **dSource** staging directories are in the form of "guid-staging-xx", where xx is the staging database number
 - b. **VDB** directories are appended with "guid-vdb-xx".
2. When all **VDBs** and **dSources** are disabled and you are ready to move the **DelphixConnector** location, stop the **DelphixConnector** service.
3. Backup the remaining directories as a precaution, in particular, the **logs** directory.
4. Uninstall the **DelphixConnector**, using the instructions from [Uninstall the Delphix Connector](#).
5. Reinstall the **DelphixConnector** to the new location, such as "C:\Delphix\DelphixConnector". Check that the **DelphixConnector** service has started.
6. Modify the new **DelphixConnector** location.
 - a. On Delphix engine versions **prior to 5.1.2.0**, use the CLI to modify the directory
 - i.

```
de4350.dcenter host> select winhost.delphix.com
de4350.dcenter host 'winhost.delphix.com'> update
de4350.dcenter host 'winhost.delphix.com' update *> set
toolkitPath="C:\Delphix\DelphixConnector"

de4350.dcenter host 'winhost.delphix.com' update *> commit
Dispatched job JOB-3203
HOST_UPDATE job started for "winhost.delphix.com".
HOST_UPDATE job for "winhost.delphix.com" completed
successfully.
```
 - ii. Refresh the windows environment for the change to take effect.
 - b. On engine versions **5.1.2.0 and higher**, you only need to refresh the windows environment.
7. Enable the **dSources** and/or **VDBs**.

Related Topics

- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Replacing Self-signed Certificates on the Delphix Connector](#)
- [Uninstalling the Delphix Connector Service from the Target Database Servers](#)

Replacing Self-signed Certificates on the Delphix Connector

The Delphix Connector relies on a Java Keystore with a self-signed X.509 certificate in order to instantiate SSL. If this certificate does not conform to the customer's business standards, it is possible to run a PowerShell script (ReplaceConnectorKeystore.ps1) to replace the self-signed certificate with a certificate that is signed by a Certificate Authority of their choice (i.e. Verisign).

This script should only be used to replace the self-signed certificate in the Delphix Connector's Java Keystore with a signed certificate. Upon execution, the script will do the following:

1. Validate that a PrivateKeyEntry exists within the input keystore
2. Stop the DelphixConnector service
3. Rename the existing DelphixConnector keystore
4. Import the new keystore
5. Start the DelphixConnector service

Prerequisites:

1. The Delphix Connector is installed
2. The DelphixConnector.jks file exists at <Drive>:\<path to DelphixConnector>\connector\DelphixConnector.jks
3. The DelphixConnector.properties file exists at <Drive>:\<path to DelphixConnector>\connector\DelphixConnector.properties and has not been tampered with (STOREPASS, KEYPASS, UUID are present)
4. The Java Keytool utility exists at <Drive>:\<path to DelphixConnector>\jre\bin\keytool.exe
5. The script, ReplaceConnectorKeystore.ps1 exists at <Drive>:\<path to DelphixConnector>\connector\ReplaceConnectorKeystore.ps1

User Inputs:

1. A JKS/PKCS#12 formatted keystore containing a PrivateKeyEntry with a signed certificate
2. The alias of the PrivateKeyEntry in the new keystore
3. The password for the new JKS/PKCS#12 keystore
4. The password for the private key in the new JKS/PKCS#12 keystore

Running the script:

Open up a PowerShell console, and do the following:

1. Navigate to where ReplaceConnectorKeystore.ps1 lives
2. Run .\ReplaceConnectorKeystore.ps1
3. Enter the full path to the new JKS/PKCS#12 keystore
4. Enter the alias of the PrivateKeyEntry in the input keystore
5. Enter the password for the input keystore
6. Enter the password for the private key in the input keystore

How to check if your Java Keystore contains a PrivateKeyEntry:

```
PS C:\Program Files\Delphix\DelphixConnector\jre\bin> .\keytool.exe -list
-keystore ..\..\connector\DelphixConnector.jks
-storepass <STOREPASS from DelphixConnector.properties file>
Keystore type: JKS
Keystore provider: SUN
Your keystore contains 1 entry
delphixconnector-4ef488a8-85df-4418-b56d-1e61b25c0aa2, Jul 28, 2017, PrivateKeyEntry,
```

```
Certificate fingerprint (SHA1): 67:79:DA:E2:64:7A:74:42:62:CA:13:66:29:16:81:0A:B9:7E:4A:60
```

Example of a successful keystore replacement:

```
PS C:\Users\dtully\Documents> .\ReplaceConnectorKeyStore.ps1
Enter the full path to a JKS/PKCS#12 keystore: C:\Program
Files\Delphix\DelphixConnector\jre\bin\test.jks
Enter alias: leaf
Enter keystore password: *****
Enter private key password: *****
Verifying that a PrivateKeyEntry exists in C:\Program
Files\Delphix\DelphixConnector\jre\bin\test.jks
Stopping the Delphix Connector service
Renaming C:\Program Files\Delphix\DelphixConnector\connector\DelphixConnector.jks to
C:\Program Files\Delphix\DelphixConnector\connector\DelphixConnector.jks.old
Importing the keystore into DelphixConnector.jks
[Storing C:\Program Files\Delphix\DelphixConnector\connector\DelphixConnector.jks]
Starting the Delphix Connector service
```

Related Topics

- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Relocating the Delphix Connector](#)

Uninstalling the Delphix Connector Service from the Target Database Servers

On occasion, the installed Delphix Connector may become out of date as newer versions of Delphix are released to improve stability or performance.

The following steps can be used to uninstall and reinstall the Delphix Connector as required:

- Uninstall the Delphix Connector
- Download the latest installer from your Delphix Engine
- Reinstall the Delphix Connector

While the Delphix Connector is uninstalled on target/staging environments, dSource and virtual database (VDB) operations (such as validated sync, manual snapshots, and provisioning operations) initiated by the Delphix Engine against those environments will fail.

When working with target/staging environments, or Delphix Engines with a large number of users, to reduce the risk of confusion or failure for other users consider disabling each dSource in the environment. In the case of VDBs, there is no need to disable or shutdown. If SnapSync, refresh or other VDB operations occur during the uninstall/install they can be restarted once the procedure is complete. This prevents the need to incur more downtime.

Uninstalling the Delphix Connector

Before uninstalling the Delphix Connector, the following steps should be performed using the Delphix Management application:

1. Login to the Delphix Management application as an admin.

 As a precaution, you can backup the binary directories, such as bin, connector, service, scripts, etc, to a temporary location.

2. Copy the logs directory, since uninstall will remove this directory. VDB or Staging folders can not be moved since the data is an iSCSI data mount. These will not be removed or manipulated during uninstall/reinstall.
3. Verify that no VDB provisioning or refresh operations are currently being performed on the target/staging environments.
4. If the environment being upgraded is a staging environment for any dSources, disable all dSources using this environment.
5. Just prior to disabling dSources, use the following query on each SQL Server instance to verify that no operations are currently running.

Show Restore Operations

```
SELECT r.session_id, r.command, CONVERT(NUMERIC(6,2), r.percent_complete) AS
[PercentComplete], CONVERT(VARCHAR(20), DATEADD(ms,r.estimated_completion_time,
GetDate()),20) AS [ETACompletionTime], CONVERT(NUMERIC(10,2),
r.total_elapsed_time/1000.0/60.0) AS [ElapsedMin], CONVERT(NUMERIC(10,2),
r.estimated_completion_time/1000.0/60.0) AS [ETAMin], CONVERT(NUMERIC(10,2),
r.estimated_completion_time/1000.0/60.0/60.0) AS [ETAHours], CONVERT(VARCHAR(10
00), (SELECT SUBSTRING(text, r.statement_start_offset/2, CASE WHEN
r.statement_end_offset = -1 THEN 1000 ELSE (r.statement_end_offset-
r.statement_start_offset)/2 END) FROM sys.dm_exec_sql_text(sql_handle)))
FROM sys.dm_exec_requests r
WHERE command IN ('RESTOREDATABASE', 'BACKUPDATABASE', 'RESTORELOG', 'BACKUPLOG')
```

 If your Windows Server is used as a Staging / Validated Sync Target for one or more linked dSources, uninstalling the connector while a Validated Sync operation (i.e. database restore) is in progress may leave the database in an inconsistent state. Validated Sync operations are not visible in the Delphix GUI, so it is necessary to query the SQL Server instance(s) on this server directly to check whether a restore is in progress.

6. If any operations are in progress, wait for them to complete before disabling dSources. Once complete, disable dSources and proceed with the uninstall.

On the Target/Staging Environment (where the Delphix Connector is being upgraded):

1. From the Windows Services manager, Stop the **Delphix Connector Service**.
2. Via the **Task Manager** verify that the java process associated with the **Delphix Connector** has terminated.
3. From **Control Panel** select **Programs and Features**, then navigate to **Delphix Connector**, right-click it and select **uninstall**.
4. Install the new connector from the downloaded **DelphixConnectorInstaller.msi** or **DelphixConnectorInstaller.exe** file. For more information refer to [Installing the Delphix Connector Service on the Target Database Servers](#).
5. Check the **Delphix Connector** version from the **Programs and Features** list. For more information on Connector versions please refer to [How To Determine Connector Version](#).
6. Check the **Delphix Connector** service is running and set to automatic.
7. Verify your system is now working properly. You may need to manually start any disabled dSources or VDBs from the Delphix UI. For more information refer to [Managing Data Sources and Syncing Data](#).

Related Links

- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Relocating the Delphix Connector](#)
- [Replacing Self-signed Certificates on the Delphix Connector](#)
- [Troubleshooting Delphix Connector](#)

Troubleshooting Delphix Connector

Troubleshooting Removing the Delphix Connector

If there are issues removing the original Delphix Connector please refer to these steps:

- You can remove the DelphixConnector from Control Panel > Program and Features.

Troubleshooting Delphix Connector Install

If you experience the following issues when re-installing the Delphix Connector refer to, [Reinstall/Upgrade the Delphix Connector](#).

- Delphix Connector service might disappear from the service console.
- Installation failing with unidentified service error.
- The Delphix connector upgrade failed.
- Manually deleted the Delphix Connector folder to try reinstalling.

Troubleshooting if the Delphix Connector Service Fails to Start

- Check if the Delphix Connector service is available and running on the Staging Host. If it is not running, then start the service.
- If the Delphix Connector service automatically shuts down or has restarting problems, kill the running java process started by Delphix connector and then retry starting the Delphix Connector Service.
- Open a command prompt (cmd) as administrator.
- From the prompt run: `netstat -noba | findstr "9100 java"`. If the connector is still running, the output should be similar to this:

```
C:\Users\username>netstat -noba | findstr "9100 java"
[java.exe]
  TCP    0.0.0.0:9100          0.0.0.0:0           LISTENING
1928
[java.exe]
  TCP    [::]:9100           [::]:0              LISTENING
1928
[java.exe]
```

The final number on each line is the Process ID (PID) of the Delphix Connector process.

- If there is a java process bound to port 9100 as shown above, but the Delphix Connector still fails to start, you can kill the Java process using the taskkill command:

```
C:\Users\username>taskkill /PID 1928 /F
SUCCESS: The process with PID 1928 has been terminated.
```

Alternatively, you can use Task Manager or Process Explorer to identify and end the correct process/PID and end it from there.

- Start the Delphix Connector service.

If the service still fails to start after performing the above steps, please attempt to reinstall the Delphix Connector. If issues persist, contact Delphix Support for assistance.

Troubleshooting if Installation of Third Party Software Results in Delphix Connect Service to Fail to Start

- If third party software is installed on the Windows host running Delphix Connector and upon the reboot, Delphix Connector doesn't startup, please uninstall the Software and revert the host back to the state it was in before the software deployment and then check if the Delphix Connector service starts
- If it doesn't start check the steps in the section "The Delphix Connector service fails to start".
- If the service still doesn't start please open a Support case with Delphix for further investigation.

Related Links

- [Installing the Delphix Connector Service on the Target Database Servers](#)
- [Relocating the Delphix Connector](#)
- [Replacing Self-signed Certificates on the Delphix Connector](#)

Manual Discovery for SQL Server Instances

Occasionally, SQL Server instances cannot be automatically discovered due to a problem in the Windows registry or if the instance is not running at that time. These SQL Server instances can be discovered using manual discovery.

Manually Adding a SQL Server Instance

To manually add a SQL Server instance

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments** and select the **Databases** tab.
4. Click the **Add Instance** button and enter the following details:
 - a. To allow linking and/or provisioning, select **Enabled**.
 2. If this environment will be used as a staging environment, select **Enabled**.
 3. Selecting **Enabled** for Fulltext Installed allows users to do full-text searches.
 4. Add as a **Failover Cluster Source**: When this is Enabled users are asked to enter the Server name. **Note**: This option is only applicable when adding a SQL Server failover cluster instance as a source, it is not available for clusters as a target.
 - e. **Name**: Not currently used.
 - f. **Port**: Used for IPC connection to the instance.
 - g. **Installation path and instance owner**: Installation path of SQL Server. Not currently used.
 - h. **Internal version**: The internal version of the SQL Server. This is important during linking and provisioning as the compatible instance is identified based on the internal version.
 - i. **Version**: the SQL Server Version required for adding a source instance.
Note: If the wrong version is entered, Delphix may fail to discover the databases and hence manual discovery may fail.
- Click **Add**. At this time, if all values are correct, an instance will be created and all its databases will be auto-discovered internally.

 If Delphix Engine is not able to establish a connection using given details, appropriate errors will be displayed and the user will be asked to make the required changes.

Editing a SQL Server Instance

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, select the **Databases** tab.
5. Click the **Pencil** icon to edit an instance.
If you are editing an instance which you added, you will be able to edit:
 - You will be able to edit all fields except the name.
6. Click **Save**.

Refresh a manually added SQL Server Instance

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, click the **refresh** icon on top.
5. Along with auto-discovered instances, it will now also refresh manually added instances and add/remove its databases, based on the current state of the instance.

 If Delphix Engine is not able to establish a connection to the instance, then the appropriate warning will be displayed to the user during database discovery.

Deleting a manually added SQL Server Instance

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, select the **Databases** tab.
5. You will see a delete icon at the far right of all manually added instances.
6. Click **Delete**. This will delete the instance and all its databases.

 This operation will fail if the chosen instance has databases that have dependencies such as database being used to provision VDBs etc.

Related Topics

- [Environment Attributes for Hosts with SQL Server](#)

Adding a SQL Server Source Environment

This topic describes how to add a SQL Server source environment.

Prerequisites

- You must have already set up SQL Server target environments, as described in [Adding a SQL Server Standalone Target Environment](#)
 - You will need to specify a target environment that will act as a proxy for running SQL Server instance and database discovery on the source, as explained in [Overview of Setting Up SQL Server Environments](#)
- Make sure your source environment meets the requirements

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions** menu and select **Add Environment**.
5. In the **Add Environment** wizard, Host and Server tab select:
 - a. Host OS: **Windows**
 - b. Host Type: **Source**
 - c. Server Type:
 - If you are adding a Windows Server Failover Cluster (WSFC), add the environment based on which WSFC feature the source databases use:
 - Failover Cluster Instances
Add the environment as a **standalone** source using the **cluster name** or **address**.
 - AlwaysOn Availability Groups
Add the environment as a **cluster** source using the **cluster name** or **address**.
 - Otherwise, add the environment as a **standalone** source.
6. Click **Next**.
7. In the Environment Settings tab select a **Connector Environment**.
Connector environments are used as proxy for running discovery on the source. If no connector environments are available for selection, you will need to set them up as described in [Adding a SQL Server Standalone Target Environment](#). Connector environments must:
 - have the Delphix Connector installed
 - be registered with the Delphix Engine from the host machine where they are located.
8. Enter the **Environment Name**, **Node Address**, **OS Username**, and **OS Password** for the source environment.
9. Click **Submit**.

As the new environment is added, you will see multiple jobs running in the Delphix Admin Job History to Create and Discover an environment. In addition, if you are adding a cluster environment, you will see jobs to Create and Discover each node in the cluster and their corresponding hosts. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel. If you don't see it, click the **Actions** menu and select **Refresh All**.

Related Topics

- [SQL Server Support and Requirements](#)
- [Adding a SQL Server Standalone Target Environment](#)

Adding a SQL Server Standalone Target Environment

This topic describes how to add a SQL Server standalone target environment to the Delphix Engine.

You can use SQL Server targets for three purposes in a Delphix Engine deployment. They can:

- Host a target environment for the provisioning of Virtual Databases (VDBs).
- Host a staging database for a linked dSource and run the validated sync process.
- Serve as a proxy host for database discovery on source hosts.

Regardless of the specific purpose, all Windows targets must have the Delphix Connector installed to enable communication between the host and the Delphix Engine. The instructions in this topic cover initiating the Add Target process in the Delphix Engine interface, running the Delphix Connector installer on the target machine, and then verifying that the target has been added in the Delphix Engine interface.

 When target environments are discovered, Delphix will configure the Microsoft iSCSI Initiator Service for Automatic startup.

Prerequisites

- Make sure that your target environment meets the requirements described in the following sections:
 - [SQL Server Support and Requirements](#)
 - [Requirements for Windows iSCSI Configuration](#)
 - [Receive Side Scaling \(RSS\) for Windows Staging Target and Targets](#)
- The Directory on which you install the Delphix Connector should have at least 1GB of available space.

 If you have taken a snapshot on Windows 2012 or an earlier version, the provisioning, linking, or exporting to Windows 2022 will result in disk errors. Windows event logs list these errors. You can ignore these errors or run the `CHKDSK /F` command.

Procedure

1. From the machine that you want to use as a target, start a browser session and connect to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions** menu and select **Add Environment**.
5. In the **Add Environment** wizard Host and Server tab, select:
 - a. Host OS: **Windows**
 - b. Host Type: **Target**.
 - c. Server Type: **Standalone**.
6. Click **Next**.
7. In the Environment Settings tab click the download link for the **Delphix Connector Installer**. The Delphix Connector will download it to your local machine.
8. On the Windows machine that you want to use as a target, run the Delphix Connector installer. Click **Next** to advance through each of the installation wizard screens. The installer will only run on 64-bit Windows systems. 32-bit systems are not supported.
 - a. For **Connector Configuration**, make sure there is no firewall in your environment blocking traffic to the port on the target environment that the Delphix Connector service will listen to.
 - b. For **Select Installation Folder**, either accept the default folder or click **Browse** to select another.
 - c. Click **Next** on the installer final 'Confirm Installation' dialog to complete the installation process and then **Close** to exit the Delphix Connector Install Program.

9. Return to the Delphix Management application.
10. Enter the **Environment Name**, **Host Address**, **Delphix Connector Port**, **OS Username**, and **OS Password** for the target environment.
11. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
12. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
13. Click **Submit**.

As the new environment is added, you will see two jobs running in the **Delphix Admin Job History**, one to **Create and Discover** an environment, and another to **Create** an environment. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel.

Post-Requisites

1. On the target machine, in the **Windows Start Menu**, click **Services**.
2. Select **Extended Services**.
3. Ensure that the **Delphix Connector** service has a **Status** of **Started**.
4. Ensure that the **Startup Type** is **Automatic**.

Related Topics

- [SQL Server Environments and Data Sources](#)

Adding a SQL Server Failover Cluster Target Environment

This topic describes how to add a SQL Server failover cluster target environment to the Delphix Engine.

Adding a failover cluster target environment will discover SQL Server failover cluster instances that are running. You can then provision virtual databases (VDBs) to these failover cluster instances.

Prerequisites

- You must add each node in the Windows Failover Cluster individually as a standalone target environment using a non-cluster address. See [Adding a SQL Server Standalone Target Environment](#).
 - For a cluster node added as a standalone environment, the Delphix Engine will only discover **non-clustered** SQL Server instances.
 - For a cluster target environment, the Delphix Engine will only discover SQL Server **failover cluster** instances.
- Each clustered SQL Server instance must have at least one clustered disk added to the clustered instance resource group, which can be used for creating mount points to Delphix storage.**
 - The clustered drive must have a drive letter assigned to it.
 - The clustered drive must be formatted using the "GUID Partition Table (GPT)" partition style in order for the Delphix Engine to automatically discover the drive letter as a valid option for the cluster instance. An MBR-formatted disk requires manual verification outside of Delphix that the disk has been correctly added to the MSSQL clustered resource group prior to creating the VDB. When provisioning the VDB, you must manually specify the desired MBR disk, because it will not appear in the Delphix GUI.
 - The clustered drive must be added to the clustered instance resource group as a dependency in the Failover Cluster Manager.
- Each node in the cluster must have the **Failover Cluster Module for Windows Powershell** feature installed.
 - While running Windows PowerShell as an administrator, enter this command to load the module:


```
import-module failoverclusters
```
- An additional target environment that can be used as a **Connector Environment** must exist. This environment must NOT be a node in the cluster. See [Adding a SQL Server Standalone Target Environment](#).

Hotfix required for Windows 2008 R2 hosts

The following hotfix is required for Windows 2008 R2 Cluster nodes:

"0x80070490 Element Not found" error when you enumerate a cluster disk resource by using the WMI MSCluster_Disk class query in a Windows Server 2008 R2-based failover cluster

<http://support.microsoft.com/kb/2720218>

Windows Cluster Volume Management Software Requirements

Only cluster volumes managed by the native Windows Volume Manager are supported. For example, cluster volumes managed by Veritas VxVM are not supported.

If you use third-party volume management software, create a new LU (recommended to be 10GB in size) and add this LU as a clustered resource to the SQL Server instance using native Windows volume management tools. Assign a drive letter for this LU. You can then use this LU as the volume mount point location for Delphix VDB provisioning.

Best Practices

SQL Server failover cluster instances that the Delphix Engine will use should not be used to host databases other than Delphix VDBs.

Cluster environment restrictions

You cannot use failover cluster target environments as staging environments.

 However, for VDB **target hosts**, it is important to use the **same edition** of SQL Server software as the **source database**, so that all features available in the source are also available in the VDB.

Supported Roles for each Instance Type

Failover Cluster Instances, and instances supporting Always On Availability Groups, support a subset of the operations available to Standalone SQL Server instances.

Color		Supported?		
		Yes		
		No		
		Environment Role		
Instance Type	Added As	Source Environment	Staging Environment	Target Environment
Standalone Instance	Standalone	 *		
Failover Cluster Instance (FCI)	Standalone	 *		
Failover Cluster Instance (FCI)	Cluster			
Always On Availability Group (AG)	Cluster	Supported		 **

* Databases that are participating in Availability Groups will not be discovered during the discovery of a Standalone environment.

** VDBs cannot be provisioned into availability groups. However, SQL Server instances that participate in an Availability Group can also be added as Standalone Instances.

Using a Failover Cluster Instance as both Source and Target

A Failover Cluster Instance added as an environment once (as either a Source or Target environment) cannot be used as both a Source and Target.

If this is required, the environment can be added twice:

- Once as a Standalone Source environment
- Once as a Cluster Target environment

The Standalone Source environment can be used for linking dSources, and the Cluster Target environment is used for provisioning VDBs.

As suggested by the Best Practice note earlier in this article, this is not a recommended configuration. Where possible, SQL Server failover cluster instances that the Delphix Engine will use as a target should not be used to host databases other than Delphix VDBs.

⚠ If you have taken a snapshot on Windows 2012 or an earlier version, the provisioning, linking, or exporting to Windows 2022 Failover Cluster Target might result in disk errors. In such a case, contact Delphix Customer Support

⚠ The following error may be encountered during the Windows 2022 AG and FCI clusters creation process on VMware ESXi.

```
An error occurred while executing the test.
```

```
Unable to connect to <hostname.domainname.com> via WMI. This may be due to networking issues or firewall configuration on <hostname.domainname.com>.
```

```
Invalid namespace
```


This failure is due to `Validate switch enabled Teaming configuration validation`. The workaround is to skip this validation.

Procedure

1. Click **Manage**.
2. Select **Environments**.
3. Next to **Environments**, click the **Plus** icon and select **Add Environment**.
4. In the **Add Environment** wizard, Host and Server tab select:
 - a. Host OS: **Windows**
 - b. Host Type: **Target**.
 - c. Server Type: **Cluster**.
5. Click **Next**.
6. In the Environment Settings tab specify the environment Name and cluster address of the Windows Failover Cluster.
7. Select a host that is NOT a node in the cluster as the **Connector Environment**.
8. Enter the **OS Username** and **OS Password** for the target environment.
9. Click **Validate Credentials**.
10. Click **Submit** to confirm the target environment addition request.

In the Delphix Engine interface, you will see a new icon for the target environment, and two jobs running in the **Delphix Admin Job History**: one to **Create and Discover** an environment, and another to **Create** an environment. When the jobs are complete, click the icon for the new environment, and you will see the details for the environment.

Example Environment



In this example environment, the Delphix Connector was installed on **Connector Environment**, **Cluster Node 1**, and **Cluster Node 2**. Each host was added to Delphix as standalone target environments.

Next, the **Windows Failover Cluster** was added as a Windows Target Cluster environment using the cluster address. **Cluster Node 1** is currently the active node for the SQL Server Failover Cluster resource group. Delphix has exported iSCSI LUs and has created the corresponding Cluster Disk resources for each VDB.

Related Topics

- [Overview of Setting Up SQL Server Environments](#)
- [Adding a SQL Server Standalone Target Environment](#)
- [Adding a SQL Server Source Environment](#)

Additional SQL Server Environment Topics

This section covers the following topics:

- [Microsoft SQL Server Bypass PowerShell Execution Policy](#)
- [Delphix as a Backup Solution to SQL Server](#)
- [Delphix Managed Backups on Secondary Nodes of SQL Server AlwaysOn Availability Group Cluster](#)
- [Receive Side Scaling for Windows Staging Target and Targets](#)
- [Validated Sync Environment](#)
- [Changing the Hostname \(IP Address\) or IQN of a SQL Server Target or Staging Host](#)

Microsoft SQL Server Bypass PowerShell Execution Policy

Hooks are executed with “PowerShell -ExecutionPolicy RemoteSigned” whereas user-configured pre/post sync and pre/post provision PowerShell scripts are executed with “PowerShell -ExecutionPolicy Bypass”. If this PowerShell configuration exceeds the security privileges of business standards, a Group Policy Object can be configured in Windows to manage PowerShell execution policy. If configured, we require that the script execution policy is set to “RemoteSigned”. The primary benefit is that the Group Policy configuration takes precedence and overrides the execution policy set in PowerShell, even if you run PowerShell as an Administrator. Alternatively, using MSSQL hooks instead of pre/post scripts will naturally execute with a “RemoteSigned” execution policy.

Refer to the Microsoft Documentation on Execution Policies for more information: https://docs.microsoft.com/en-us/powershell/module/microsoft.powershell.core/about/about_execution_policies?view=powershell-6

 The Group Policy is per machine and will apply to PowerShell 2.0 and later.

Related Topics

- [Additional SQL Server Environment Topics](#)

Delphix as a Backup Solution to SQL Server

Using Delphix as a Backup Solution to SQL Server

Delphix provides you the option to automatically manage backups from SQL Server source databases into the Delphix Engine. Prior to Delphix 4.2, users could not link source databases that were backed up by unsupported backup software. In Delphix 4.2, a feature known as Delphix Managed Backups was introduced that allows you to have Delphix take and manage backups from your source database directly into Delphix storage. This is the first step in Delphix becoming a full-featured backup solution for SQL Server databases. When the Delphix Engine manages the backups for a dSource, it takes regular, copy-only full backups of the source database, so activating the feature will not interfere with existing backup management solutions. You can configure the schedule of when the Delphix Engine takes these copy-only full backups by specifying a SnapSync policy for the dSource. You can change the SnapSync policy for a dSource at any time by visiting the policy screen; there, you can either select a new SnapSync policy or modify the current one.

If you use a backup solution that is not supported by Delphix, you cannot use your existing backups to keep your dSources in sync. However, enabling Delphix Managed Backups will overcome this issue by using automatic copy-only full backups to keep dSources in sync. Currently, dSources linked when this feature is enabled will not support LogSync functionality, which means that you can only provision VDBs from snapshots and not from any time between snapshots. Additionally, in the current release, the Delphix Engine cannot take differential or transaction log backups of the source database.

Linking SQL Server dSources with Delphix Managed Backups

The **Data Management** page of the link wizard for SQL Server dSources provides the option to enable Delphix Managed Backups.

It possible to enable this feature here at link time or toggle it on after the link for future syncs. If you enable this feature, the dSource can only use copy-only full backups taken by the Delphix Engine to stay in sync with its source; the Delphix Engine will prohibit syncing using existing backups. Checking the **Enabled** box results in the following changes to the **Data Management** page:

- The initial load option is set to a copy-only full backup taken by the Delphix Engine
- The ability to provide a backup path disappears
- A SnapSync policy selection screen appears in the **Policies** page

You can select from the list of existing SnapSync policies if the one you want doesn't exist you will need to create a new one on the **Policies** page under the **Manage** dropdown. Proceeding through the remainder of the link wizard will create a dSource with Delphix-managed backups enabled. You can confirm that a dSource has the feature by selecting the dSource and going to the **Configuration > Data Management** tab after creation and checking the **Delphix Managed Backups** section, as displayed below:


Delphix_Admin


Timeflow	Status	Configuration		
Source	Policies	Data Management	Masking	Hooks
VALIDATED SYNC CONFIGURATION				
Delphix Managed Backups Yes				
Force Compression No				
Encryption Key Yes				
Netbackup Ingestion Disabled				
Commvault Ingestion Disabled				

To disable/enable this feature after linking:

Toggling Delphix Managed backups is supported for SQL Server dSources. This means a dSource that was previously created with Delphix Managed Enabled can have regular validated sync and similarly a dSource that had external backups can use Delphix Managed Backups.

To disable Delphix Managed Backups for a specific dSource, select that specific dataset and go to **Data Management** tab under the **Configuration** tab.


Delphix_Admin


Timeflow
Status
Configuration

Source
Policies
Data Management
Masking
Hooks

VALIDATED SYNC CONFIGURATION

Delphix Managed Backups ⓘ

Enabled

Autodiscover Backup Path

Enabled

Encryption Key

.....

Validated Sync Mode

Transaction Log ▼

LogSync

Netbackup Ingestion ⓘ

Enabled

Commvault Ingestion ⓘ

Enabled

×
✓

Click the edit button and uncheck this feature. Then update all new inputs.

To enable this feature just go to the same edit menu.

Backup Compression Feature Overview

Since Delphix Engine 5.2, the Delphix Engine has allowed compression to be enabled ("forced") for SQL Server backups which use Delphix Managed Backups. Delphix Managed Backups are used to synchronize SQL Server dSources when existing SQL Server backup files cannot be made available to the Staging server, or if a third-party

backup vendor is used that is not yet supported by Delphix. For more information on this functionality see [Delphix as a Backup Solution to SQL Server](#) and [Linking a dSource with SQL Server](#).

Backup compression is preferable in the following situations when:

- The default backup compression setting for the Source SQL Server instance is Disabled (0)
- The Source database does not use Transparent Data Encryption (TDE)
- The Source server has available CPU resources to perform compression

Where the compression ratio of a compressed backup exceeds 3:1, this will reduce the amount of data (by half) that must be transferred over the network to perform a backup, resulting in much faster SnapSync operations.

How to Enable Backup Compression

If the SQL Server instance's default backup compression setting is enabled (see Microsoft's document [View or Configure the backup compression default Server Configuration Option](#)), no specific action is required. Backups will automatically be compressed in accordance with this setting.

If backups are not automatically compressed, the Delphix Engine (5.2 and later) can be configured to force compressed backups using the following steps:

- Login to the Delphix Engine's Management interface.
- Open the **Manage > Datasets** screen.
- Locate and select the dSource (using the tree on the left).
- Navigate to the **Configuration > Data Management** pane.
- Click the **Edit** (pencil) icon to modify the Validated Sync settings.

AdventureWorks2014

Timeflow

Status

Configuration

Source	Policies	Data Management	Masking	Hooks
--------	----------	------------------------	---------	-------

Validated Sync Configuration


Delphix Managed Backups

Yes



- Select the **Force Compression** checkbox, then the select the Tick icon to save the changes

Validated Sync Configuration



The change will automatically take effect during the next scheduled Snapshot/SnapSync operation.

Delphix Managed Backups on Secondary Nodes of AlwaysOn Availability Group Cluster

The Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012 (11.x), Always On availability groups maximizes the availability of a set of user databases for an enterprise. An *availability group* supports a failover environment for a discrete set of user databases, known as *availability databases*, that failover together. An availability group supports a set of read-write primary databases and one to eight sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations. Some customers prefer using secondary nodes for backup to reduce the load on the primary node while some customers are OK with using the primary node. To provide more flexibility for customers using Always On availability groups with Delphix Managed backups, delphix provides an option to choose backup policy during linking of a dSource.

They can pick one of the following options that describe their backup policy as shown below.

- Primary (**default option**): Backups are taken only on the primary node.
- Secondary Only: Customers who never want to use primarily for their backups must choose this option as it ensures backups never go to the primary node. If the secondary nodes are down, the backups fail but do not use primary at all.
- Secondary preferred: As the name suggests, one of the secondary nodes is used for backup. The backup will be taken on primary if none of the secondary nodes can be used for backup. So, customers who want reliable backups even if they are taken on the primary node if required should use this option.

Data Management
Configure and Administer data details.

Managed Backups ⓘ

Enabled

Initial Load
Delphix will take a copy-only full backup of your source database

Force Compression

Backup Policy

Primary only
Backups only go to the primary node.

Secondary only
Backups only go to secondary nodes. If secondary nodes are down, backups will fail.

Secondary preferred
Backups go to secondary nodes, but if secondary nodes are down, backups will go to the primary node.

The backup policy can be changed after dSource creation too as shown below.

Timeflow Status **Configuration**

Source Policies **Data Management** Masking Hooks

Validated Sync Configuration

Delphix Managed Backups ⓘ

Enabled

Force Compression

Backup Policy


Primary only
Backups only go to the primary node.

Secondary only
Backups only go to secondary nodes. If secondary nodes are down, backups will fail.

Secondary preferred
Backups go to secondary nodes, but if secondary nodes are down, backups will go to the primary node.

Encryption Key


The backup policy is not only provided during the linking phase of a dsource but also during the manual snapsync phase too to provide more flexibility to customers as shown below. Default option is the Backup Policy selected during linking phase.



Understanding SnapSync Policies

SnapSync policies provide you the ability to specify how frequently the Delphix Engine takes a copy-only full backup of a source database when Delphix Managed Backups are enabled. Selecting an initial SnapSync policy is mandatory at dSource link time. However, you can change the SnapSync policy that the Delphix Engine applies to a dSource at any time by visiting the policy management screen:

1. Click **Manage**.
2. Click **Policies**.



Check SnapSync Policy

For dSources that have Delphix-managed backups enabled, the current SnapSync policy will be displayed under the **SnapSync** column. The rows corresponding to dSources that do not use Delphix Managed Backups will be grayed out. Clicking the **current SnapSync policy** for a dSource will display a drop-down menu of existing SnapSync policies along with the option to create a new SnapSync policy. Selecting a SnapSync policy from this list will change the current SnapSync policy for the dSource. When creating a new policy, you will see the following screen:

Create New SnapSync Policy

Here, you can configure the frequency with which the Delphix Engine takes backups of your source database. You can modify these schedules at any time by clicking the **Modify Policy Templates** button in the upper right-hand corner of the policy management screen.

The **Timeout** field above specifies how long a SnapSync job is allowed to run before it is terminated. If a SnapSync job exceeds its timeout window, the Delphix Engine discards the new backup and rolls back the dSource to the most recent snapshot.

Related Topics

- [Additional SQL Server Environment Topics](#)

Delphix Managed Backups on Secondary Nodes of SQL Server AlwaysOn Availability Group Cluster

The Always On availability groups feature is a high-availability and disaster-recovery solution that provides an enterprise-level alternative to database mirroring. Introduced in SQL Server 2012 (11.x), Always On availability groups maximizes the availability of a set of user databases for an enterprise. An *availability group* supports a failover environment for a discrete set of user databases, known as *availability databases*, that fail over together. An availability group supports a set of read-write primary databases and one to eight sets of corresponding secondary databases. Optionally, secondary databases can be made available for read-only access and/or some backup operations. Some customers prefer using secondary nodes for backup to reduce load on the primary node while some customers are OK with using primary node. To provide more flexibility for customers using Always On availability groups with Delphix Managed backups, delphix provides an option to choose backup policy during linking of dSource.

They can pick one of the following options that describe their backup policy as shown below.

- **Primary (default option):** Backups are taken only on the primary node.
- **Secondary Only:** Customers who never want to use primarily for their backups must choose this option as it ensures backups never go to the primary node. If the secondary nodes are down, the backups fail but do not use primary at all.
- **Secondary preferred:** As the name suggests, one of the secondary nodes is used for backup. The backup will be taken on primary if none of the secondary nodes can be used for backup. So, customers who want reliable backups even if they are taken on the primary node if required should use this option.

Data Management
Configure and Administer data details.

Managed Backups ⓘ

Enabled

Initial Load
Delphix will take a copy-only full backup of your source database

Force Compression


Backup Policy

Primary only
Backups only go to the primary node.


Secondary only
Backups only go to secondary nodes. If secondary nodes are down, backups will fail.

Secondary preferred
Backups go to secondary nodes, but if secondary nodes are down, backups will go to the primary node.

The backup policy can be changed after dSource creation too as shown below.



The backup policy is not only provided during the linking phase of a dsource but also during the manual snapsync phase too to provide more flexibility to customers as shown below. The default option is the Backup Policy selected during the linking phase.



Related Topics

- [Additional SQL Server Environment Topics](#)
- [Delphix as a Backup Solution to SQL Server](#)

Receive Side Scaling for Windows Staging Target and Targets

Enabling Receive Side Scaling (RSS) on a Windows Target and Staging Target can have a significant improvement in the overall IO throughput to the Delphix Engine and is a best practice. RSS enables network adapters to distribute the kernel-mode network processing load across multiple processor cores in multi-core computers. The distribution of this processing makes it possible to support higher network traffic loads than would be possible if only a single core were to be used.


More information on RSS can be found [here](#).

⚠ Enabling RSS on the network interface will force the network service to restart and will cause a momentary loss of connectivity on that network interface.


i Because hyper-threaded CPUs on the same core processor share the same execution engine, the effect is not the same as having multiple core processors. For this reason, RSS does not use hyper-threaded processors.

Steps to implement RSS on Windows


1. From Server Manager/Local Server/Network Connections select the NIC that Delphix will be connecting to.



2. Select Properties and then Configure.



3. From the Property menu on the left, select **Receive Side Scaling**, confirm that it is set to **Enabled**, and select **OK** to close each of the open windows.



Related Topics

- [Additional SQL Server Environment Topics](#)

Validated Sync Environment

This topic describes additional requirements for SQL Server environments that will be used as targets for validated sync. You must configure a staging (Validated Sync) environment as a target, with a few additional requirements.

Requirements for SQL Server Validated Sync Target Environments

Each SQL Server target environment used for validated sync must meet these requirements:

- Only standalone target environments can be used as validated sync target environments. Windows Failover Cluster target environments and SQL Server Failover Cluster instances cannot be used.
- The SQL Server instance must be the same version as the instance hosting the source database. For more information about compatibility between different versions of SQL Server, see [SQL Server Support Matrix](#)
- The owner of the SQL Server instances on the target environment that are used for the staging databases must have SMB read access to the location containing the backup images of the source databases
- If the source database is backed up with third-party backup software like LiteSpeed or Red Gate SQL Backup Pro, you must install the backup software on both the source and the validated sync environment. For backup software compatibility requirements, see [SQL Server Support Matrix](#)

Add the Validated Sync Environment

The order is important. Add the validated sync environment as the first step in setting up the SQL Server topology.

1. From the machine that you want to use as a target, start a browser session and connect to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions** menu and select **Add Environment**.
5. In the **Add Environment** wizard Host and Server tab, select:
 - a. Host OS: **Windows**
 - b. Host Type: **Target**.
 - c. Server Type: **Standalone**.
6. Click **Next**.
7. In the Environment Settings tab click the download link for the **Delphix Connector Installer**. The Delphix Connector will download to your local machine.
8. On the Windows machine that you want to use as a target, run the Delphix Connector installer. Click **Next** to advance through each of the installation wizard screens. The installer will only run on 64-bit Windows systems. 32-bit systems are not supported.
 - a. For **Connector Configuration**, make sure there is no firewall in your environment blocking traffic to the port on the target environment that the Delphix Connector service will listen to.
 - b. For **Select Installation Folder**, either accept the default folder, or click **Browse** to select another.
 - c. Click **Next** on the installer final 'Confirm Installation' dialog to complete the installation process and then **Close** to exit the Delphix Connector Install Program.
9. Return to the Delphix Management application.
10. Enter the **Environment Name, Host Address, Delphix Connector Port, OS Username, and OS Password** for the target environment.
11. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
12. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
13. Click **Submit**.

As the new environment is added, you will see two jobs running in the **Delphix Admin Job History**, one to **Create and Discover** an environment, and another to **Create** an environment. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel.

Related Topics

- [Additional SQL Server Environment Topics](#)

Changing the Hostname (IP Address) or IQN of a SQL Server Target or Staging Host

This topic describes how to change the hostname (IP address) or iSCSI Qualified Name (IQN) of a Windows Target or Staging host.

By default, Windows servers generate an IQN based on the hostname assigned to it. Changing the hostname will change the host IQN as well. Because the Delphix Engine exports storage for dSources and VDBs to Windows hosts using iSCSI, changes to the Windows hostname must be made according to the following procedure. If you have set a non-default IQN on a Windows Target or Staging host, and want to change that IQN, you must follow these procedures.

 Changing the hostname or IQN of a Windows target or staging server requires that you modify the iSCSI Initiator configuration on the Windows host. Doing so incorrectly can cause failures in dSources, VDBs, or non-Delphix users of iSCSI on the Windows host. The instructions in this topic describe how to change the IQN using the `iscsicli` command-line utility. Because many people are less familiar with the `iscsicli` utility, the instructions also include information for using the iSCSI Initiator graphical user interface. Failing to carefully follow the steps below in sequence can cause availability issues for your dSources and VDBs. If you have questions about the following instructions, [please contact Delphix Support for help](#).

1. Disable the dSources.
2. Disable the VDBs.

 If your Windows server has dSources or VDBs from more than one Delphix Engine, you will need to disable the dSources and VDBs on each Delphix Engine.

3. Remove any remaining persistent volumes from the Windows server. From the Server Manager\Tools\iSCSI Initiator configuration tool, use the options available in the **Volumes and Devices** tab.
or,
Follow these steps to use the `iscsicli` command-line utility:
 - a. List the persistent volumes.

```
PS C:\> iscsicli reportpersistentdevices
Microsoft iSCSI Initiator Version 6.1 Build 7601
Persistent Volumes
"\\?\storage#volume#{bb38add1-d03f-11e1-8767-005056b37fe6}
#0000000008010000#{53f5630d-b6bf-11d0-94f2-00a0c91efb8b}"
"C:\Program Files\Delphix\DelphixConnector\564d6fbb-df9d-e90b-00f1-
da37b17011d3-staging-15\ARCHIVE\"
[...]
```

The operation completed successfully.

- b. Volumes with a "normal" path correspond to mounted volumes. For example, `C:\Program Files\Delphix` is a normal path. If you see any normal paths in the output, be sure you have disabled all of the VDBs and dSources.
- c. Volumes with a path beginning `\\?\` correspond to unmounted persistent volumes. Remove each of them:

```
PS C:\> iscsicli RemovePersistentDevice
"\\?\storage#volume#{bb38add1-d03f-11e1-8767-005056b37fe6}
#0000000008010000#{53f5630d-b6bf-11d0-94f2-00a0c91efb8b}"
```

- d. Alternatively, if all of the persistent devices are for unmounted volumes, you can remove them all at once with this command:

```
PS C:\> iscsicli clearpersistentdevices
```

4. Remove all of the persistent targets.

From the Server Manager\Tools\iSCSI Initiator configuration tool, use the options available in the **Favorite Targets** tab.

or,

Follow these steps to use iscsicli command-line utility:

1. List persistent targets:

```
PS C:\> iscsicli ListPersistentTargets
```

2. Remove the appropriate persistent targets. Below is sample output listing the persistent targets:

```
PS C:\> iscsicli ListPersistentTargets
[...]
Target Name           : iqn.2008-07.com.delphix:02:02843619-12c4-e4d2-8041-
f5c56a647bc2
Address and Socket    : 10.43.5.45 3260
Session Type         : Data
Initiator Name       : Root\ISCSIPRT\0000_0
Port Number          : <Any Port>
Security Flags       : 0x0
Version              : 0
Information Specified: 0x20
Login Flags          : 0x0
Username             :
[...]
```

Misleading Help for RemovePersistentTarget Command

The help for iscsicli `RemovePersistentTarget` is misleading:

```
iscsicli RemovePersistentTarget <Initiator Name> <TargetName>
                                <Port Number>
                                <Target Portal Address>
                                <Target Portal Socket>
```

<Initiator Name> and <Target Name> show up in the listing and should be taken directly from there.

<Port Number> can be taken from the listing output, but a * should be used if <Any Port>

is listed.

<Target Portal Address> and <Target Portal Socket> are shortened to Address and Socket in the ListPersistentTargets output. The term Socket in both places is what is more typically referred to as a port .

- Use the RemovePersistentTarget command to remove the target, as shown in this example:

```
PS C:\> iscsicli RemovePersistentTarget Root\ISCSIPRT\0000_0 iqn.2008-07.com.delphix:02:02843619-12c4-e4d2-8041-f5c56a647bc2 * 10.43.5.45 3260
```

- Log out of any sessions.
From the Server Manager\Tools\iSCSI Initiator configuration tool, use the options available in the **Targets** tab to log out. Selected a connected session under **Discovered Targets**, and then click **Disconnect**.

or,

Follow these steps to use the iscsicli command-line utility:

- List the sessions.

```
PS C:\> iscsicli sessionlist
Session Id : fffffa8003fb0018-4000013700000001
Initiator Node Name : iqn.1991-05.com.microsoft:10-43-1-200.ad.delphix.com
Target Node Name : (null)
Target Name : iqn.2008-07.com.delphix:02:02843619-12c4-e4d2-8041-f5c56a647bc2
[...]
```

- Log out from the target.

```
PS C:\> iscsicli logouttarget fffffa8003fb0018-4000013700000001
```

- Change the hostname (IP address) or IQN in the Delphix engine.
 - If you are changing the hostname, follow the instructions in the [Microsoft TechNet](#) article "Rename the Computer."

 Note that if the computer is on a domain, you will need a domain administrator to perform the rename or re-add the computer to the domain depending on the version of Windows it is running.

- To use the Delphix Setup application, go to **Manage** → **Environments** screen and select the Target host. To modify the Host Address for the environment, use the Edit (Pencil) icon next to the **Attributes** panel. Update the Host Address to reflect the new IP address, and then use the **Save (Tick)** icon to confirm the change. This will automatically trigger an action to Refresh the Environment. The refresh process includes steps to configure the iSCSI Target and verify that it is reachable.
 - If you are changing the IQN only, change it through the Microsoft iSCSI Initiator GUI following the instructions in the [Microsoft iSCSI User Guide](#).
- Wait for the computer to finish rebooting.

- Verify the new IQN in the iSCSI initiator.

 If you are using the default IQN and have changed the hostname (IP address), the IQN should include the new hostname.

- Refresh the environment on the Delphix Engine.
- Re-enable the dSources.
- Re-enable the VDBs.
- Using the `iscsi cli` command-line utility, verify that the sessions on the Windows server are using the new IQN.

```
PS C:\> iscsicli sessionlist
Microsoft iSCSI Initiator Version 6.1 Build 7601

Total of 1 sessions

Session Id : fffffa8003f77018-4000013700000004
Initiator Node Name : <NEW IQN>
[...]
```

Related Topics

- [Microsoft TechNet article "Renaming the Computer"](#)
- [Microsoft iSCSI User Guide](#) (download)

Linking Data Sources and Syncing Data with SQL Server

Creating a dSource will ingest data from the source and create a dSource on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

This section covers the following topics:

- [Linking a dSource from a SQL Server: An Overview](#)
- [Linking Data Sources with SQL Server](#)
- [Staging Push Implementation for SQL Server](#)
- [Data Management Settings for SQL Server Data Sources](#)
- [Additional SQL Server dSource Topics](#)

Linking a dSource from a SQL Server: An Overview

When linking a dSource from a SQL Server source database, Delphix offers several different methods of capturing backup information:

- Delphix Managed Backups, where the Delphix Engine schedules and initiates the backups from the source database, and captures them
- SQL Server Managed Backups, where the SQL Server source database schedules and initiates backups and the Delphix Engine captures them
 - Full backups
 - Full or differential backups
 - Transaction log backups (with LogSync disabled)
 - Transaction log backups (with LogSync enabled)

Delphix Managed Backups are conceptually a lot simpler to explain, but they can be considered less desirable because they might be viewed as intrusive. SQL Server Managed Backups are explained in two sections:

- the initial load of the dSource from the source database
- subsequently keeping the dSource synchronized with the source database.

Below is a brief explanation of how these three different modes of operation work.

Delphix Managed Backup Mode

When the **Enabled** checkbox for Delphix Managed Backups is selected, the Delphix Engine will initiate a COPY_ONLY full backup of the source database for the initial load of the dSource. Thereafter, the Delphix Engine will initiate COPY_ONLY full backups of the source database using the schedule specified by the selected SnapSync Policy. If SnapSync Policy is set to None, the Delphix Engine will not automatically initiate COPY_ONLY backups, but you can initiate them manually using the snapshot (camera) icon.

When the **Force Compression** checkbox for Delphix Managed Backups is selected and the backups are not compressed, the Delphix Engine will take a compressed copy-only full backup of the source database and this will take effect during the next scheduled Snapshot or SnapSync operation.

 COPY_ONLY backup files are written to the Delphix storage that has been mounted on the staging server. Delphix does not require space on the source or staging servers to hold the backup files.

You must also select the Staging Environment and the SQL Server instance onto which the backups will be restored.

SQL Server Managed Backup Modes

Initial Load of the dSource

When the Delphix Managed Backups option is left unchecked (*default*), the Delphix Engine will initiate a backup only when the user selects to initiate a COPY_ONLY full backup that the Delphix Engine will use to keep the dSource in sync with the source database.

For the initial load of the dSource, you can choose one of the following:

- have the Delphix Engine initiate a COPY_ONLY full backup
- use the most-recent existing full or differential backup (*default*)
- use a specific existing full or differential backup identified by its `backup_set_uuid`

 **Simple Recovery Model**

If the source database is using a simple recovery model, using a new full COPY_ONLY backup initiated by the Delphix Engine is not supported for the initial load of a dSource.

After the initial load, you need to select the Backup Paths and tell the Delphix Engine where to look for backups of the source database.

- If **Autodiscover** is selected, the exact path used to take the backup from the source database will be determined by querying by the source database instance. If this option is used, the Delphix Engine should take backups to a UNC path (Windows file share) so that they are accessible to the Staging Server.
- If custom paths are specified, the Delphix Engine will query the source database instance, identify the filename of the source backup, and then recursively search the specified Backup Paths for this filename. These paths should also be UNC paths (Windows file share) which are accessible to the Staging Server.

The path used to restore backups must be readable by the Windows server hosting the staging instance, using the staging environment's configured Environment User.

The Delphix Engine supports source database backups that SQL Server creates natively, as well as backups created by Quest/Netvault LiteSpeed, Red Gate SQL Backup Pro, Veritas NetBackup, and Commvault. For more information, see the topic [SQL Server Support Matrix](#).

Once you have decided how the dSource will be initially loaded, select the staging environment and the SQL Server instance onto which the backups will be restored. NetBackup and Commvault backups are not on local storage and therefore the Backup Paths will just be ignored. See [Linking a dSource from a NetBackup SQL Server Backup](#) and [Linking a dSource from a Commvault SQL Server Backup](#) for more information.

 The staging instance opens the backup file for reading and may hold a lock on it when restoring the backup on the staging database. A new source database backup, initiated with the **Append to the existing backup set** option, may fail as SQL Server will not be able to open the locked backup file to append a new backup to it.

Keeping the dSource Synchronized with the Source Database

Next, specify how the Delphix Engine will capture subsequent backups of the source database.

The selected Validated Sync mode determines how often the Delphix Engine will check for new backups, and which type of backups it will check for. You can always force synchronization with the source database by enabling **Validated Sync Mode** from the **Data Management** tab available under the **Configuration** tab for the selected dSource.

Validated Sync and LogSync

When you link a source database into the Delphix Engine, a staging database will still be required if the Validated Sync is not enabled, as described in [Overview of Setting Up SQL Server Environments](#). In this process, the Delphix Engine continuously monitors the source database for new full and differential backups if the source database is using a **simple** recovery model, or transaction log backups if using a **full** recovery model. This will also depend on the selected backup mode. When it detects a new backup, it restores that backup to the staging database with the storage residing in Delphix. The result is a TimeFlow with consistent points from which you can provision a VDB, also known as snapshots.

Snapshots accumulate over time. To view a snapshot:

1. From the **Datasets** panel, click the **group** containing the dSource.
2. Select **dSource**.
3. Click the **TimeFlow** tab.

Each snapshot is displayed and includes some information about the captured database along with Snapshot database change number (SCN for Oracle and LSN for SQL Server). You can scroll through these cards to select the one you want, or you can enter a date and time to search for a specific snapshot.

Summary of Validated Sync Modes

This table summarizes each mode of Validated Sync, displaying how often the Delphix Engine will poll to check for new backup files when it creates snapshots for the dSource, and whether point-in-time restores for provisioning and refreshing virtual databases (VDBs) is possible or not.

Validated Sync Mode	Polling Interval	Snapshot for each FULL backup	Snapshot for each DIFF backup	Snapshot for each TLOG backup	Allows Point-in-time Restores	Notes
Delphix Managed Backups	N/A	✗	✗	✗	✗	Takes COPY_ONLY full backups according to SnapSync schedule. For more information on this option, see Delphix as a Backup Solution to SQL Server .
Transaction log backups (LogSync DISABLED)	1-minute	✗	✗	✓	✗	Log backups are not collected if: <ul style="list-style-type: none"> • There are gaps in the sequence of log backups (a break in the "log chain") • The available log backups do not include any changes since the last successful Delphix snapshot
Transaction log backups (LogSync ENABLED)	1-minute	✗	✗	✓	✓	Log files consume additional space on the Delphix Engine and are managed according to the defined retention policy for logs. For NetBackup and Commvault backups, Point-in-time restores are not supported.

Full or differential backups	1-minute	✓	✓	✗	✗	
Full backups	1-minute	✓	✗	✗	✗	
None	Manual only	✗	✗	✗	✗	Only retrieves backups when you initiate a manual snapshot.

Timeflow cards

The Delphix Engine will create a Timeflow card for each backup it restores to the staging server. For example:

- A database in Full backups Validated Sync mode, and daily backup configured on the dSource, would receive one Timeflow card per day
- A database in Full or differential backups sync mode, with one daily backup and two differential backups per day, would receive three Timeflow cards per day
- A database in Transaction log backups sync mode, with a log backup every 15 minutes, would receive 96 TimeFlow cards per day

Related Topics

- [SQL Server Support and Requirements](#)
- [Delphix as a Backup Solution to SQL Server](#)
- [Linking a dSource from a Commvault SQL Server Backup](#)
- [Linking a dSource from a NetBackup SQL Server Backup](#)

Linking Data Sources with SQL Server

Linking a dSource will ingest data from the source and create a dSource object on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

For an overview of all dSource related actions, please [Managing Data Sources and Syncing Data](#).

When linking a dSource from a SQL Server source database, Delphix offers several different methods of capturing backup information:

- SQL Server Managed Backups, where the SQL Server source database schedules and initiates backups and the Delphix Engine captures them
 - Full backups
 - Full or differential backups
 - Transaction log backups (with LogSync disabled)
 - Transaction log backups (with LogSync enabled)
- Delphix Managed Backups, where the Delphix Engine schedules and initiates the backups from the source database, and captures them

Procedure

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.
4. In the **Add dSource** wizard, select the source database with the correct environment user-specified.
5. Select user type for source database authentication and enter the login credentials. Enter username and password for Database user or Domain (Windows) user. For Environment User, select a source environment user from the dropdown list and click **Next**.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed. For more information, [Data Management Settings for SQL Server Data Sources](#).
8. Select the Staging environment and SQL Instance that will be used to manage the staging database used for validated sync of the dSource.
9. Select any policies for the new dSource.
10. Enter any scripts that should be run on the **Hooks** page.
11. Review the dSource Configuration and Data Management information, and then click **Submit**.

Related Topics

- [Data Management Settings for SQL Server Data Sources](#)
- [Additional SQL Server dSource Topics](#)

Staging Push Implementation for SQL Server

This topic provides implementation details of Staging Push for SQL Server. Staging Push eliminates the need of accessing the customer's production environment and hence the dependency on it. For end-users with a wide variety of architectural requirements such as an unsupported backup appliance, complex ingestion requirements, and supporting alternate replication methods, Staging Push increases flexibility by ensuring that Delphix can work with most of these unique requirements while maintaining Delphix standard product support.

Staging Push gives end-users control over some Staging DB processes so that the nuanced, staging-based ingestion can be orchestrated externally. It will give control of the staging database to the end-user to pull from any backup provider (as a part of this, you'll be responsible for keeping the Staging DB in sync with the production database). Staging database files will be stored on Delphix Storage. Delphix will still be the one snapshotting the underlying data files, and gathering the metadata required to provision from the snapshot.

The below steps show how to create a dSource using the Staging Push mechanism.

Procedure

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.
4. On the Preparation tab, click **Next**.
5. From the **dSource Type** tab, select the **Staging Push** option and click **Next**.
6. From the **dSource Configuration** tab, enter the following dSource configurations, also shown in the screenshot below, and click **Next**.

dSource Name - This name will be available on the Delphix Engine interface.

Database Name - This is the name of the staging database that is created on the staging host after linking. It is suggested to use the word "staging" while assigning a name to the staging database as this will help to

distinguish it from other end-user databases.

dSource Configuration

Target a group from your Datalist or create a new group

dSource Name

Database Name ⓘ

Target Group Add Dataset Group

Untitled
▼

Notes

7. From the **Data Management** tab, select the staging environment and repository details. The staging environment is where the staging database will be hosted and a *repository* is a container for the *SourceConfigs* objects. Each environment in Delphix can contain any number of SQL Server instances and each SQL Server instance can contain any number of databases.
8. Select any policy for the new dSource. SnapSync policy is used as a default policy for taking snapshots. For more details on SnapSync policy, see [Policies for Scheduled Jobs](#).
9. Enter any script that should be run on the **Hooks** page. For staging push dSources, a restore backup script can be added as a part of the hooks script to automate backup restore. Refer to the below section for more details.
10. Review the Staging Push dSource Configuration and Data Management information, and then click **Submit**.

Restoring Backups

After linking and before taking a snapshot, if the user wants to align the staging database with the source database automatically they can do so by providing the restore script in the hook, as shown below.

They can also manually restore a database backup on the staging host.

Also, refer to [SQL Server Hook Operation Notes](#) for environment variables used with staging push dSources.

⚠ First backup restoration needs to be full backup with REPLACE keyword.

Prerequisites

⚠ When restoring the backup, make sure that all the database files are present at the mounting location.

The mounting location is available as **Mount Base** under the **Configuration** tab of the newly created dSource.

The screenshot shows the 'Dsource' configuration interface. The 'Configuration' tab is selected, and the 'Source' sub-tab is active. The configuration is divided into two main sections: 'STAGING DATABASE' and 'STAGING ENVIRONMENT'.

STAGING DATABASE	STAGING ENVIRONMENT
Name Dsource1	Name sg-win2019tgt
Size 200.00MB	OS Windows (Windows)
Version MS Sql 13.2.5233.0	Timezone America/Los_Angeles,PST-0800
Mount Base C:\Program Files\Delphix\DelphixConnector\564d18cc-f930-62d6-3bff-b86b9dbc959f-staging-8	Repository SQL2016
Recovery Model FULL	User qa-ad\delphix

Procedure

Manually restore a backup on the staging host or provide the restore script in the hook script. A sample pre-sync hook for native backup is shown below. It also lists the following keywords.

- NORECOVERY - Keeps the database in restoring state.
- REPLACE - Overwrites existing database created during linking with whichever source database that is in the backup set and is getting restored.
- MOVE - Restores the data and log file to the specified locations.
- BACKUP_FILE_LOCATION - The location of the backup file of the source database from where the backup is restored and it should be accessible from the staging host. For example, \\10-43-89-18\Backup\sp.bak.

Sample hook to illustrate restore script usage for restoring backups

```
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
#Set-Variable UTF8_CODEPAGE 65001 -option readonly
```

```

#Set-Variable UTF8_ENCODING "System.Text.UTF8Encoding" -option readonly

# Uncomment the following log to turn on debugging
# Set-PSDebug -Trace 2;

$DSOURCE_HOST = $env:STAGING_INSTANCE_HOST
$DSOURCE_PORT = $env:STAGING_INSTANCE_PORT
$DSOURCE_INSTANCE = $env:STAGING_INSTANCE_NAME
$CONNECT_STRING = "$DSOURCE_HOST\$DSOURCE_INSTANCE,$DSOURCE_PORT"
$DSOURCE_NAME = "$env:STAGING_DATABASE_NAME"
$DataDbFilePath = "$env:STAGING_MOUNT_BASE"
$SQL_SCRIPT= "RESTORE DATABASE $DSOURCE_NAME FROM DISK = '\
\BACKUP_FILE_LOCATION' WITH NORECOVERY, REPLACE, MOVE 'sp' TO
'$DataDbFilePath\DATA\db\$DSOURCE_NAME.mdf', MOVE 'sp_log' TO
'$DataDbFilePath\DATA\db\' + $DSOURCE_NAME + "_log.ldf'"

function die {
    Write-Error "Error: $($args[0])"

    # run exit handler, if defined
    if (Get-Command -type Function -name atExit 2> $null) {
        atExit
    }
    exit 1
}

function verifySuccess {
    if (!$?) {
        die "$($args[0])"
    }
}

### Restore database.

echo $SQL_SCRIPT
Sqlcmd -b -S $CONNECT_STRING -U sa -P <PWD> -Q $SQL_SCRIPT
verifySuccess "Failed to restore backup"

# Uncomment the following line to turn off debugging
# Set-PSDebug -Trace 0;

exit 0

```

Snapshot

The first snapshot is created as a part of dSource creation and contains data and log files within the Delphix created mount point.

 The first snapshot created is of an empty database and does not contain any source database's data. (Unless a source backup was restored via a pre-sync hook. In that case, the initial snapshot will be that of source backup and won't be empty).

Prerequisites

- The database should be present on the staging host.
- The DATA directory should be mounted.
- The DB files should be present on Delphix mounted DB directory. For example, `C:\Program Files\Delphix\DelphixConnector\ec2197b2-e0c6-48d2-bd14-265e6fa9b5ab-staging-1\DATA\db`
- The database should be in Restoring state.
- No other restore operation should be in progress on the staging database.

SnapSync Criteria

- Delphix tries to fetch the last restored backup. If no backup is found, the snapshot is skipped with a warning.
- If a snapshot already exists for the backup in the current timeflow, the snapshot is skipped with a warning.

 Snapshots display the Staging Host timezone, as opposed to Linked dSources, where snapshots display the Source Host timezone. Here, the Staging timezone is displayed for Staging Push as we don't have the Source host to fetch the timezone information. This functioning might change in the future.

Perform the following steps to take a snapshot:

1. Login to the **Delphix Management** application.
 2. Click **Manage** and select **Datasets** from the dropdown list.
 3. Select the dSource created using the Staging Push mechanism to Snapshot.
 4. Click the **Camera** icon.
 5. From the Snapshot dialog box, select **Yes**.
 6. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the Snapshot you just created.
- You can now proceed to provision the VDB using the snapshot.

Disabling and Enabling the dSource

When a disabled dSource is later enabled, it will resume communication and incremental data updates from the staging database according to the original policies and data management configurations that you set.

Procedure

Disabling a dSource will stop further operations on the Delphix Engine related to the staging dSource.

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **dSource** you want to disable.
5. In the upper right-hand corner, from the **Actions** menu (...) select **Disable**.
6. In the Disable dialog select **Disable**.

When you are ready to enable the dSource again, from the Actions menu (...) select **Enable**, and the dSource will continue to function as it did previously.

Keep a note of the following:

- During enable, if there exists a database with the same staging database name as the user entered and it is not Delphix managed, or if it is but Delphix Engine is unable to drop it, the enable operation will not be terminated.

- Storage mounting will be attempted for the dSource and a new job event `SKIP_ATTACH_OPERATION` will be displayed. This means the job is updated with a warning and attach operation is skipped. The enable operation will be marked as successful.

 Customers can drop the database that is outside Delphix or that wasn't dropped and try again the disable/enable option.

- Delphix will attempt to mount the staging database's storage during the enable job. Even if the attach job fails, the storage will remain mounted (if Delphix was able to successfully mount it) and the enable job will be marked as successful.

Unlink(Detach)/Link(Attach) a dSource

Each dSource contains metadata that associates it with the staging database, as well as the data it has ingested from the staging database in the form of snapshots up to that point. It is possible to detach, or unlink, a dSource from its staging database.

- A detached dSources can still be used to provision a virtual database (VDB).
- You can re-link the staging push dSource with a different staging database name than before. In that case, the staging database will be created with the new name provided. However, DB file names will remain the same as before.
- Delphix Engine supports converting existing Linked dSource to Staging push dSource and vice versa.

Unlinking or Detaching a dSource

1. Login to the **Delphix Admin** application.
2. Click **Manage**.
3. Select **My Datasets**.
4. Select the **database** you want to unlink.
5. From the Actions ... tab, click **Unlink**.
6. Click **Yes** to confirm.

Attaching a Previously Detached dSource

1. Login to the **Delphix Admin** application.
2. Click **Manage**.
3. Select **My Datasets**.
4. Select the **database** you want to link.
5. From the Actions ... tab, click **Link dSource**.
Select **Staging Push** and enter **staging database Name**, select **Staging Environment** and **Staging Repository**.

Link dSource ✕

dSource Type

Linked dSource
Staging Linked to a source database

Staging Push
Customer managed staging database

Staging Database Name

Staging Environment

Staging Repository

Cancel
Link

6. Click **Link** to confirm.

Limitations

1. Restore backups fails on staging database due to mount issues.
When a critical threshold is enforced due to Delphix storage and later resumed successfully, a transactional log will not be applicable on the staging database directly. A generic warning will be displayed for the resume job only for staging push dSources on the Delphix interface that after the mount, the user might face issues in restoring a transaction log.
SQL server error when restoring the Tlog after resume:

```
Msg 4320, Level 16, State 3, Line 13
The file " SourceDB4_log " was not fully restored by a database or file restore. The entire file must be
successfully restored before applying this backup set.
Msg 3013, Level 16, State 1, Line 13
RESTORE LOG is terminating abnormally.
RESOLUTION: Full or differential backup will be successfully restored.
```
2. Attach operation fails during the Enable operation.

If the attach operation fails for staging push dSource, any exception is consumed and the enable operation is marked as a success.

For example, if DB files were changed for the staging database but a snapshot was not taken on the Delphix engine and disable was performed directly, the enable will succeed with an attach failure that will be shown as a warning. The previous DB files will be available on the mount path.

RESOLUTION: The customer can apply a full restore at this point to bring the database in restoring mode before trying to run a sync.

3. Point-in-Time provisioning is not supported.
4. The sub-directory structure for DB files is not supported.

For example,

If a staging database is restored with DB files inside another folder within the DB directory (`C:\Program`

`Files\Delphix\DelphixConnector\ec2197b2-e0c6-48d2-bd14-265e6fa9b5ab-staging-1\DATA\db\folder1\sp.mdf`, `C:\Program`

`Files\Delphix\DelphixConnector\ec2197b2-e0c6-48d2-bd14-265e6fa9b5ab-staging-1\DATA\db\folder1\folder2\sp_log.ldf`), then operations such as to enable,

relink, VDB provisioning, and export will fail.

5. The validated sync process is not supported for Staging Push dSources.
6. Staging database states other than Restoring such as ReadOnly, StandBy, and Online are not supported. If a database is in any of these unsupported states, dSource will go into the Cannot Monitor state on the Delphix Engine interface, and Sync operation will fail.

Data Management Settings for SQL Server Data Sources

Each dSource has its own data management settings, which can be configured during the linking workflow as well as in the configuration page for that dSource.

You can configure data management settings to improve overall performance and match the needs of your specific server and data environment.

The following settings are available for SQL Server data sources:

Setting	Explanation
Managed Backups	When enabled, the Delphix engine will take full backups of the database, per the dSource's SnapSync policy, and validated sync will be disabled. Existing backups cannot be used to synchronize the dSource when backups are managed by Delphix.
Recovery Model	The current recovery model of the source database. Three recovery models exist simple, full, and bulk-logged. This is not configurable from Delphix but can be changed within the source's database settings.
Initial Load	<ol style="list-style-type: none"> 1. Delphix will take a copy-only full backup of your source database. 2. Use the most recent full or differential backup (default). 3. Use a specific full or differential backup.
Backup Paths	<p>These are the locations where Delphix will be looking for backups for ingestion.</p> <p>Select Autodiscover to have the Delphix Engine automatically locate the backups by querying the msdb database in the SQL instance.</p> <p>Otherwise, for each path, please specify the top level of a particular backup path since the Delphix Engine will recursively search for the backup file in all subdirectories beneath the specified path.</p>
Validated Sync Mode	<p>Determines the types of backups validated sync will use to generate snapshots.</p> <ol style="list-style-type: none"> 1. Transaction log backups. 1. LogSync adds log files from the source database to the dSource, allowing you to provision a virtual database (VDB) from a specific point in time or LSN for SQL Server databases. <ul style="list-style-type: none"> • Full or Differential backups. • Full backups. • None.
Staging Environment	This environment will host the staging database used for validated sync.
Repository	A repository is a container for the SourceConfigs objects. Each environment can contain any number of repositories, and repositories can contain any number of source configurations. A repository typically corresponds to a database installation. Whereas source configurations typically correspond to the databases.

Encryption Key	The encryption key to be used when restoring encrypted backups. If the source database is backed up using LiteSpeed or RedGate with password-protected encryption, you can supply the encryption key that the Delphix Engine should use to restore those backups.
NetBackup Ingestion	Enables ingestion from a NetBackup source.
Commvault Ingestion	Enables ingestion from a Commvault source.

Related Topics

- [Linking Data Sources with SQL Server](#)
- [Additional SQL Server dSource Topics](#)

Additional SQL Server dSource Topics

This section covers the following topics:

- [Changing the Staging Target Environment for a SQL Server dSource](#)
- [Upgrading a dSource after a SQL Server Upgrade](#)
- [Linking a dSource from a Commvault SQL Server Backup](#)
- [Linking a dSource from a NetBackup SQL Server Backup](#)
- [Restoring SQL Backups Stored in Azure Cloud Storage](#)
- [Working with SQL Server Snapshots](#)
- [Detaching and Re-Attaching SQL Server dSources](#)

Changing the Staging Target Environment for a SQL Server dSource

This topic describes how to change the staging target environment for a SQL Server dSource.

Prerequisites

The dSource for the staging database must be disabled before you can change the staging target environment. Follow the steps in *Enabling and Disabling SQL Server dSources* in [Managing Data Sources and Syncing Data](#) to disable the dSource.

Procedure

1. Click **Manage**.
2. Select **Datasets**.
3. Select the **dSource** for which you want to change the staging target environment.
4. Click the **Configuration** tab to view the **Staging Environment**.
5. Click the **Pencil** icon next to **Staging Environment**.
6. Edit the target server and the SQL Server instance on the server to use for staging.
7. Click the **Check** icon to save your changes.

Related Topics

- [Managing Data Sources and Syncing Data](#)
- [Linking Data Sources with SQL Server](#)
- [Data Management Settings for SQL Server Data Sources](#)

Upgrading a dSource after a SQL Server Upgrade

This topic describes how to upgrade dSources after a SQL Server database upgrade.

There are two ways to upgrade a Source database:

- Perform an Upgrade installation of SQL Server, upgrading the Source and Staging SQL Server instances in-place
- Perform a fresh installation of the new SQL Server version, and migrating the databases to the new instance

The steps required to support this in the Delphix Engine are different depending on method chosen. The required steps are outlined in the sections below.

In-place SQL Server upgrade

1. **Disable** all dSources on the instance being upgraded
2. Following Microsoft's procedures, perform an upgrade of the Source SQL Server instance
3. Following Microsoft's procedures, perform an upgrade of the Staging SQL Server instance to the same version as the Source
4. Refresh the Source and Staging environments in the GUI
5. **Enable** the dSources being upgraded

Migrate databases to newer instance

Prerequisites for database migration

- Record the configuration data for each dSource being upgraded, including the Database User, Database Password (if applicable), and Validated Sync configuration. This will be needed to re-link the dSource.
- Following Microsoft's procedures, install a new Source SQL Server instance with the new SQL Server version.
- Ensure that a Staging SQL Server instance is available, running the same SQL Server version as the new Source SQL Server instance.
- **Add** or **Refresh** the Environment containing the new SQL Server instance(s), using the steps in [Adding a SQL Server Source Environment](#) and [Adding a SQL Server Standalone Target Environment](#).

 The refresh/rediscover operations do not affect the operations of any dSources or VDBs on the environment.

Migration steps

1. Navigate to the **Manage** → **Datasets** screen.
2. For each dSource being upgraded:
 - a. Select **Unlink dSource** from the **Actions menu (...)**
3. Following Microsoft's procedures, migrate the Source database to the new Source SQL Server instance.
4. Commence a new **Full** database backup of each upgraded Source database
5. From the **Manage** → **Environments** screen, **Refresh** the new Source Environment. The migrated databases should be detected and visible from the Environment's **Databases** tab.
6. Navigate to the **Manage** → **Datasets** screen.
7. For each dSource that was upgraded:
 - a. Select **Link dSource** from the **Actions menu (...)**
 - b. Locate the upgraded Source database using the Source Environment, **Installation and Database** drop-down boxes
 - c. Select a compatible Staging Environment and Staging Repository
 - d. Configure the **Database Authentication** using values that are appropriate for the new Source database
 - e. Click **Link** to begin linking the dSource

- f. Reconfigure any dSource settings using the dataset's **Configuration** tab
- g. Use the **Snapshot** button to take a snapshot using the upgraded database backup

Related Topics

- [Additional SQL Server dSource Topics](#)

Linking a dSource from a Commvault SQL Server Backup

Customer requirements:

- Delphix currently supports Commvault v11. The version of Commvault SQL Agent on the staging environment must be the same as that on the source.
- The TCP port 8415 must be open from the staging host to the Delphix Engine.
- If the dSource is backed up with Commvault, the source and staging environments must each have the SQL Agent installed.
- Both SQL Agents (on source and staging) should be registered with Commvault Server mentioned during linking.
- The install path's 'Base' directory of Commvault SQL Server Agent on the staging host, must be part of the PATH environment variable as we need to access the Commvault CLI. This is typically located at <Commvault install path>\Base.
 - a. There are two PATH environment variables. One is for the current logged in user and the other is a global System variable. Since the Delphix Connector runs under the "Network Service" account and spawns the Commvault commands, changes need to be done to the System variable as well.
 - b. After the Commvault binaries are added to both the system and user PATH environment variables, the Delphix Connector service must be restarted in order for the new process to reflect the changes made to the PATH environment variable.
 - c. After making the changes to the PATH, login as the Delphix operating system user and try running "qlogin -sso -gt" to confirm the user can execute Commvault commands and authenticate to the Commvault server.
- Since SSO is used from Commvault CLI to login to the CommServe server from the staging environment, SSO should be enabled.
- Active Directory domain should be configured and configured staging environment user should be given required permissions to restore the database on staging client.
- In the CommServe server, the staging client should be configured with a user who has required permissions to restore the database on the staging environment. User account configuration can be done as explained in Commvault documentation at User Account Configuration for the SQL Server Agent.
- Only transaction logs, incremental, and database full backups are currently supported.

Linking with Commvault Backups:

To link a dSource and use Commvault, follow the steps to add the environments as earlier, and make sure all the requirements listed above are met. Going through the linking wizard on the **Data Management** page, select **Show advanced**.

Add dSource

- Source
- dSource Configuration
- Data Management**
- Policies
- Hooks
- Summary

Use the most recent full or differential backup

Use a specific full or differential backup

Backup Paths

Autodiscover ✎

There are no paths added.

Validated Sync Mode ⓘ

Full or Differential backups

Full backups

None

Transaction log backups

LogSync ⓘ

Enabled

Staging Environment ⓘ

virtual-tgt.dlpxdc.co ▼

Repository

SQL2005 ▼

Show advanced ▼

Enable **Commvault Ingestion** and input the CommServe hostname and, source and staging client names.

Hide advanced ^

Encryption Key ⓘ

Netbackup Ingestion ⓘ

Enabled

Commvault Ingestion ⓘ

Enabled

We validate whether the source and staging clients exist in the CommServe server and the Commvault SQL Agent on staging client is registered to the CommServe server. To run the validation commands following login command with SSO option is used to login into CommServe Server:


```
qllogin -sso -gt
```

LogSync for SQL Server dSources

Logsync (point-in-time provisioning) is currently not supported for Commvault transaction logs. However, LogSync can still be enabled if Commvault ingestion is enabled. LogSync for backups taken with other backup providers that support LogSync will work as before.

Enabling Commvault for previously created dSources

On an already created dSource go to **Configure > Data Management** and click edit (pencil button). Enable Commvault Ingestion and add all the configuration as necessary.



Linking an Availability Group database with Commvault

Linking with an Availability Group (AG) source works similar to as described above, however, the source client name provided in the Commvault configuration should be the MSSQL AG Client name as in the CommServe server.

General Notes/Troubleshooting

- Commvault backups are not on local storage and therefore while ingesting backups taken using Commvault, Backup Paths provided in configuration will be ignored.

Related Topic

- [Additional SQL Server dSource Topics](#)

Linking a dSource from a NetBackup SQL Server Backup

Customer Requirements

- The version of NetBackup client on the staging environment must be the same as that on the source.
- The TCP port 8415 must be open from the staging host to the Delphix Engine
- Backups must be taken via an MS-SQL-Server type policy with an INSTANCE client list type. (value is 15, and 1 respectively) Only transaction logs, incrementals, and database full backups are currently supported.
- The master server and source client servers' clocks must be within a minute of each other (timezones can be different).
- Any backups that Delphix needs to ingest for a dSource must be taken to one NetBackup master server using one NetBackup SQL Server client. A multiple master server setup for one dSource is not supported.
- If the dSource is backed up with NetBackup, the source and staging environments must each have the NetBackup client installed.
- Both clients (on source and staging) during installation must be setup with the master server and the source and staging instances must be registered.
- The install path's bin directory of the SQL Server NetBackup client on the staging host must be part of the system PATH as we need access to dbbackex.exe and bplist.exe. This is typically located at <client install path>\NetBackup\bin.
- If you have modified the PATH, then please restart the Delphix connector service otherwise the linking process won't be able to pick the changed system PATH environment variable.
- Configure redirected restores on the master server between the source NetBackup Client and the staging NetBackup Client
- https://www.veritas.com/support/en_US/doc/17221771-126559330-0/v113535700-126559330

Linking with NetBackup Backups

To link a dSource and use NetBackup, follow the steps to add the environments as before, and make sure all requirements listed above are met. Going through the linking wizard on the **Data Management** page, select **Show advanced**.

Add dSource

- Source
- dSource Configuration
- Data Management**
- Policies
- Hooks
- Summary

Use the most recent full or differential backup

Use a specific full or differential backup

Backup Paths

Autodiscover

There are no paths added.

Validated Sync Mode ⓘ

Full or Differential backups

Full backups

None

Transaction log backups

LogSync ⓘ

Enabled

Staging Environment ⓘ

virtual-tgt.dlpxdc.co

Repository

SQL2005

Show advanced

Enable **Netbackup Ingestion** and input the master and source client names.

Hide advanced ^

Encryption Key ⓘ

Netbackup Ingestion ⓘ

Enabled

Commvault Ingestion ⓘ

Enabled

We use these names to query bplist using these options:

```
bplist -S <master name> -C <client name> -t 15 -R -s <start date> -e <end date> \
```

If validation does not work, make sure this command actually returns backups that were taken for this dSource.

⚠ When validating the master and client servers we expect that at least one NetBackup MS-SQL-Server type backup had been taken with this pair in the last two days. Validation will fail otherwise.

Using NetBackup Config Parameters or Templates

When restoring we create a batch file based on information we find on the backup. Batch files can be customized with non-default and additional options using config templates.


https://www.veritas.com/content/support/en_US/doc/123947690-126579517-0/id-SF930853806-126579517 explains which parameters in the batch file can be edited outside of these blacklisted keywords:

- ALTCLIENT
- BROWSECLIENT
- DATABASE
- DUMPOPTION
- ENABLESERVICEBROKER
- ENDOPER
- MOVE
- NBIMAGE
- NBSERVER
- OBJECTNAME
- OBJECTTYPE
- RECOVERED STATE
- RESTOREBEFOREMARK
- RESTOREBEFOREMARK AFTERTIME
- RESTOREOPTION
- RESTOREPAGES
- PARTIAL

- PAGE

Any key that is used outside of what is documented may cause the restore to fail with unknown errors.

If the **Configure NetBackup Config Template** is checked, the page after **Data Management** becomes **NetBackup Config Template**.



This page allows this dSource’s specific config parameters to be updated, either by adding rows to the table or inputting them as text input (key=value).

⚠ For config params/templates, the 'value' will be injected into the script as is so if the value of the parameter needs to be a String, quotes should be included with 'value'.

If creating a new set of config parameters while linking each dSource seems unnecessary, a config template can be specified. However, that is only possible via the CLI. To do that, log onto the CLI using the Delphix admin user/ password and go to **database > templates**. Make sure to create the new template with sourceType set to MSSqlLinkedSource.

Example:

Then when linking, select the template created. (In this screenshot it would be **NetBackup**) If any edits are made to a selected template while linking, this will create a new set of config parameters for the specific dSource and will not edit the actual template. Config parameters added to a specific dSource will be ignored if the dSource has a config template selected.

LogSync for SQL Server dSources

LogSync can still be enabled if NetBackup ingestion is enabled however Point-in-Time provisioning is currently not supported for NetBackup transaction logs. LogSync for backups taken with other backup providers that support LogSync will work as before.

Enabling NetBackup for previously created dSources

On an already created dSource go to **Configure > Data Management** and click edit (pencil button). Enable NetBackup Ingestion and add all the configuration as necessary.

Netbackup Ingestion ⓘ

Enabled

NetBackup Config Template

Configure NetBackup Config Template

Master Name

Source Client Name

Validate

If **NetBackup Config Template** is set to **Default** then the params will be empty. Go to the CLI to update the config params.

```

Password:
ip-10-110-238-136> cd source
ip-10-110-238-136 source> select Biscuit
ip-10-110-238-136 source 'Biscuit'> update
ip-10-110-238-136 source 'Biscuit' update *> edit mssqlNetbackupConfig
ip-10-110-238-136 source 'Biscuit' update mssqlNetbackupConfig *> set configParams.BLOCK_SIZE=6
ip-10-110-238-136 source 'Biscuit' update mssqlNetbackupConfig *> ls
Properties
  type: MSSqlNetbackupConfig
  configParams:
    BLOCK_SIZE: 6 (*)
  configTemplate: (unset)
  masterName: nbu-80-master.delphix.com
  sourceClientName: ln-win2012src.dlpxdc.co
ip-10-110-238-136 source 'Biscuit' update mssqlNetbackupConfig *> commit

```

Like above, to create a config template, create the template first so that during this update that specific template can be selected in the drop-down.

Delphix Managed Backups must be disabled before NetBackup can be enabled.

To disable Delphix Managed Backups for a specific dSource, select that specific dataset and go to **Data Management** tab under the **Configuration** tab.

 **Delphix_Admin** 

Timeflow	Status	Configuration		
Source	Policies	Data Management	Masking	Hooks

VALIDATED SYNC CONFIGURATION 

Delphix Managed Backups
Yes

Force Compression
No

Encryption Key
Yes

Netbackup Ingestion
Disabled

Commvault Ingestion
Disabled

Click the edit button and uncheck this feature. Then update all new inputs.

The screenshot shows the Delphix Admin interface with a navigation menu at the top. The menu includes 'Timeflow', 'Status', 'Configuration', 'Source', 'Policies', 'Data Management', 'Masking', and 'Hooks'. The 'Configuration' section is active, and a dialog box titled 'VALIDATED SYNC CONFIGURATION' is displayed. The dialog box contains the following settings:

- Delphix Managed Backups** ⓘ
 - Enabled
- Autodiscover Backup Path**
 - Enabled
- Encryption Key**
 -
- Validated Sync Mode**
 - Transaction Log ▼
- LogSync
- Netbackup Ingestion** ⓘ
 - Enabled
- Commvault Ingestion** ⓘ
 - Enabled

At the bottom right of the dialog box, there is a close button (X) and a blue confirmation button (✓).

Linking an Availability Group database with NetBackup

Linking with an AG source works similarly however the source client name inputted to the NetBackup config should be the windows cluster name rather than the client name that takes the backup.

General Notes/Troubleshooting

- If a restore fails and the staging environment's NetBackup client job log shows no attempted restores, it is likely that you need to log onto the environment using the staging user provided to Delphix. This article explains that the user profile is created on the first login and that NetBackup requires it to restore. https://www.veritas.com/support/en_US/article.100032299
- Through observation, we've noticed that the block size of the restore batch file must be equal to or less than what was used when backing up the dump. We default to the blocksize to what the backup was taken with but the config param/template can include BLOCK_SIZE to change. We won't prevent this but it's up to the user to set it correctly. Editing this field is generally not recommended for this reason.
- NetBackup backups are not on local storage and therefore the Backup Paths will just be ignored.

Related Topic

- [Additional SQL Server dSource Topics](#)

Restoring SQL Backups Stored in Azure Cloud Storage

Need for Supporting SQL Backups Stored in Azure Cloud Storage

Currently, you can restore the SQL backups that are stored locally on the staging host, or on the network path, or on the backup servers with third-party vendors such as Commvault or Netbackup. At the same time, there are users that are having their native SQL backups on Azure Cloud Storage.

Delphix now supports restoring native SQL backups from Azure Cloud Storage. This enables users who are moving to Azure to use direct backups from the Azure Storage containers instead of third-party vendors.

These backups support the following.


- This is supported by SQL Server 2016 and above
- All backup modes
 - Full backups
 - Differential
 - Transaction log backups

Logsync (point-in-time provisioning) is currently not supported. However, LogSync can still be enabled if Azure backup is being ingested. LogSync for SQL native backups that are present on Disk will work as before.
- Striped backups

All backup files are to be entirely present on the Azure Cloud and no part of a backup should be present outside of Azure Cloud.
- Validated Sync
- Azure Storage authentication mode

This backup solution uses a shared access signature token authentication method to authorize access to the blob data.

Workflow Architecture



Workflow Steps

- **User takes backup:**

You need to take the SQL native backup directly to the Azure Cloud using SQL [Backup to URL](#). Performance can be improved by enabling COMPRESSION while taking the backup. If the backup size is large, it should be striped into multiple files.

```
BACKUP DATABASE [SourceDB4]
to URL = 'https://idea1201.blob.core.windows.net/ideal1201container/source.bak'
```

- **Authentication to authorize access to the blob data:**

You must create a [SQL Credential](#) on the staging host for the following cases.

- Create using SAS authentication token as Access Key authentication is currently not supported.
- Create for each Azure container where the backup files are expected to be stored.
- Create on each SQL instance where restore has to be performed.

```
CREATE CREDENTIAL [ https://idea1201.blob.core.windows.net/
ideal1201container] (this is the Azure Container URL)
WITH IDENTITY = 'SHARED ACCESS SIGNATURE',
SECRET =
'sv=2019-12-12&ss=bfqt&srt=c&sp=rwdlacupx&se=2021-01-19T14:18:04Z&st=2021-
01-19T06:18:04Z&spr=https&sig=DBLZnu0VQTaXUwY9IgBEqNbSk';
(provide the SAS token here)
```

- **Run queries on Staging host:**

Delphix Engine uses this credential to run the following queries [Restore](#), [Restore Headeronly](#), and [Restore Filelistonly](#) on the staging host.

- Irrespective of whether auto-discovery is selected or not, if an Azure backup is being ingested, its location will be fetched from the `msdb.dbo.backupmediafamily` table on the source host and it will be restored on the staging host. Hence, it is necessary that the backup files are present on the same blob URL where the backup was originally taken. This is different from the usual native backups to the disk.

-  In SQL native disk backup,
- If auto-discovery is on, the backup path is fetched from `msdb.dbo.backupmediafamily` table on the source host.
 - If auto-discovery is off, Delphix Engine uses the custom path(s) specified by the user.
- If any native Azure backup is found, Delphix Engine will always try to restore it.
- This is similar to the user taking backups using multiple backup vendors.
 - In that case, Delphix Engine tries to restore all backups, irrespective of the backup vendor.

Unsupported Features

The following features and functionalities are currently not supported.

- Third-party backup vendors
- Point-in-time provisioning
- Access key authentication method for Azure backups
- Support for SQL Server versions below SQL 2016
- Moving backup files across Azure containers
- Moving a backup from disk to cloud
- Striped backups - Backup files that are partially present on Azure Cloud and partially on another device.
- Managed identity

Related Topics

- [Linking a dSource from a SQL Server: An Overview](#)

Working with SQL Server Snapshots

This section lists the steps to take a snapshot and delete the same.

Taking a snapshot creates a new snapshot entry in the Oracle dSource's Timeflow. You can use either **Snapshot (Default)** or **Snapshot with Parameters** option for taking the snapshot.

Snapshot (Default)

Perform the following steps to take a snapshot:

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog box, select **Yes**.
6. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the Snapshot you just created.
7. To delete the snapshot, select the snapshot you just created, and from the Actions menu (...), select **Delete Snapshot**.
8. From the **Delete Snapshot** dialog box, select **Delete**.
9. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the snapshot you just deleted.

Snapshot with Parameters

Perform the following steps to take a snapshot:

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the arrow next to the Camera icon and select the **Snapshot with Params...** option.
5. From the Snapshot dialog box, select one of the following:
 - a. **Force Full Backup** - If you select this option, then the Delphix Engine will perform an incremental backup by default. You must select this option only when a full backup is required. Full and Incremental backups consume the same space on the Delphix Engine.
 - b. **Double Sync** - Selecting this option will perform a SnapSync operation as normal. After the first SnapSync is successful, the Engine will immediately perform a second SnapSync without waiting for the Log Files required for the first SnapSync to be made consistent. This is most useful when performing the initial SnapSync (or when "Force Full Backup" is selected) on a very large database that would lead to a large number of archive logs being required to make the SnapSync consistent. Provisioning from a SnapSync that requires excessive recovery is typically time-consuming.
 - c. **Do Not Resume** - If a failure is encountered during the initial SnapSync, the Delphix Engine can resume the SnapSync at a later date. This option will cause the engine to not resume, but rather to start the initial SnapSync over again.
6. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the Snapshot you just created.
7. To delete the snapshot, select the snapshot you just created, and from the Actions menu (...), select **Delete Snapshot**.
8. From the **Delete Snapshot** dialog box, select **Delete**.
9. Navigate to the **Timeflow** tab and click **View: All snapshots** to verify the snapshot you just deleted.

Detaching and Re-Attaching SQL Server dSources

Each dSource contains metadata that associates it with the source database, as well as the data it has ingested from the source database in the form of snapshots up to that point. It is possible to detach, or unlink, a dSource from its source database. This breaks the association with the source database without affecting the data within the Delphix Engine. Detached dSources and their source databases have these properties:

- A detached dSources can still be used to provision a virtual database (VDB).
- You can re-link the source database as a different dSource.

Prerequisites

A dSource can only be attached to a new data source once it has been unlinked.

Detaching a dSource

1. Login to the **Delphix Management** application as a user with **OWNER** privileges on the dSource, group, or domain.
2. Click **Manage**.
3. Select **Databases**.
4. Select the **database** you want to unlink or delete.
5. From the **Actions** menu (...) select **Unlink**.
A warning message will appear.
6. Click **Unlink** to confirm.

Rebuilding Source Databases and Using VDBs

In situations where you want to rebuild a source database but retain the existing dSource, you will need to detach the original dSource and create a new one from the rebuilt data source.

1. Detach the dSource as described in the procedure on this page.
2. You cannot attach a dSource with the same name as a dSource that is already attached. If you intend to give the new dSource the same name as the original one, rename the detached dSource.
 - a. At the top of the **Configuration** tab, next to the dSource's name, click the **Edit** (pencil) icon.
 - b. After renaming the dSource, click the green **checkmark**.
3. Create the new dSource from the rebuilt database.

You will now be able to provision VDBs from both the detached dSource and the newly created one, but the detached dSource will only represent the state of the source database prior to being detached.

Attaching a Previously Detached dSource

You can only re-attach databases that represent the same physical database.

1. Login to the **Delphix Management** application as a user with **OWNER** privileges on the dSource, group, or domain.
2. Click **Manage**.
3. Select **Databases**.
4. Select the database you want to re-attach.
5. From the **Actions** menu (...) select **Link dSource**.
6. Select the environment and database that represents the source to which you would like to re-attach.
7. Select an appropriate staging environment, and choose the credentials that should be used to authenticate to the source database.

8. Click **Link** to link the dSource to the new source database.
9. Once linking is complete, use the **Configuration > Data Management** tab to set up an appropriate **Validated Sync** configuration.

Attaching a Previously Detached dSource using CLI

You can only re-attach databases that represent the same physical database.

1. Login to the **Delphix CLI** as `delphix_admin` or a user with OWNER privileges on the dSource, group, or domain.
2. Select the dSource by name using `database select <dSource Name>`.
3. Run the `attachSource` command.
4. Set the source config you want to attach to, using `set source.config=<Source Database Unique Name>`. Source configs are named by their database unique name.
5. Set any other source configuration operations as you would for a normal link operation.
6. Run the `commit` command.

Related Links

- [Managing SQL Server Environments and Hosts](#)

Provisioning and Managing Virtual Databases with SQL Server

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment. From a dSource, you can select a snapshot or point in time to create a VDB. SQL Server VDBs each have their own configuration settings.

This section covers the following topics:

- [Overview of SQL Server Virtual Databases](#)
- [Provisioning a SQL Server VDB](#)
- [V2P with a SQL Server VDB](#)
- [Additional SQL Server VDB Topics](#)
- [CDC Support in SQL Server](#)

Overview of SQL Server Virtual Databases

Virtual databases are a key data management concept for Delphix, explained in [Provisioning and Managing Virtual Databases](#). In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment, as described in the overview for [Managing Environments and Hosts](#) and [Overview of Requirements for SQL Server](#).

From a dSource, you can select a snapshot or point in time to create a VDB. SQL Server VDBs each have their own configuration settings as described in Configuration Settings for SQL Server Virtual Databases below. This document describes the steps to provisioning VDBs with SQL Server.

Procedure

1. In the Datasets panel on the left-hand side, click the group containing the dSource or VDB from which you want to provision and select the dSource or VDB from the provided list.
2. From the TimeFlow tab, select a snapshot to provision from. To provision from a specific point in time from dSources with LogSync enabled, use the open LogSync button.
3. Click to open the Provision VDB wizard, and select a compatible Target Environment for the new SQL Server VDB
4. On the Target Configuration page, you may customize the VDB. For a list of available configuration options, see [Configuration Settings for SQL Server Virtual Databases](#) below.
5. Select a Snapshot Policy for the VDB.
6. If the VDB should be masked during provisioning, enable Masked Provisioning by selecting an option on the Masking page
7. Enter any operations that should be run on the Hooks page.
8. Review the VDB Configuration and Summary, and then click Submit.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the Status tab, or by selecting System and viewing the Jobs page.

Alternatively, you can see this in the Actions Sidebar. When provisioning is complete, the VDB will be included in the group you designated and listed in the Datasets panel.

Configuration Settings for SQL Server Virtual Databases

Each VDB has its own data management settings, found during the provisioning workflow as well as in the configuration page for that VDB. When you create a SQL ServerVDB, Delphix copies most configuration settings from the dSource and uses them to create the VDB. However, you can customize these with the following settings:

Setting	Explanation
Recovery Model	The current recovery model of the source database. By default, this value is set to SIMPLE. You must set it explicitly during provisioning.
Auto VDB Restart	Enabling this option will automatically restart this VDB whenever its target host is rebooted.
Change Data Capture (CDC)	Indicate whether this virtual source should be enabled for CDC.

Related Topics

- [Provisioning and Managing Virtual Databases with SQL Server](#)

Provisioning a SQL Server VDB

This topic describes how to provision a virtual database (VDB) from a SQL Server dSource.

Prerequisites

- You must have already linked a dSource from a source database, as described in [Linking a SQL Server dSource](#) or have already created a VDB from which you want to provision another VDB.
- You must have already set up Windows target environments and installed the Delphix Connector on them, as described in [Adding a SQL Server Standalone Target Environment](#).
- Make sure that you have the required privileges on the target environment.
- If you are provisioning to a different target environment than the one where the staging database has been set up, make sure that the two environments have compatible operating systems. For more information on the staging database and the validated sync process.
- If using Change Data Capture (CDC):
 - The SQL Server instance on which the VDB is being provisioned or exported must support CDC.
 - The *SQL Server Agent* for the instance must be running otherwise adding a CDC capture and cleanup jobs will fail.

Procedure

 When provisioning or refreshing a SQL Server VDBs, both the DB_CHAINING and TRUSTWORTHY database parameters will be disabled (even if they were enabled on the dSource).

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource**.
5. Click **Timeflow** tab.
6. Next to a snapshot select the  Provision VDB icon.
The **Provision VDB** panel opens.
7. Select a **target environment**.
8. On the **Target Configuration** page, specify a **Mount Path**.
9. On the **Configuration** page, do the following:
 - a. Under **Target Group**, select a Target Group for the VDB.
 - b. Enter a database name.
 - c. Under Recovery Model, click the drop-down list to select a recovery model. You can choose one of the following:
 - i. Simple - This option is recommended and selected by default. It allows SQL Server to automatically mark parts of its transaction log file for re-use if they are not in use.
 - ii. Full - This option lets you take the responsibility for taking backups and log backups of the VDB to an external location.
 - iii. Bulk Logged - This option lets you take the responsibility for taking backups and log backups of the VDB to an external location.
 - d. specify any **Pre Scripts** or **Post Scripts** that should be used during the provisioning process. For CDC users, it is recommended that VDB Post Start hooks be configured to automatically start these jobs on successful provisioning.

 It is advised to use a Post Start hook as it is executed in all VDB operations like provision, refresh, rewind, disable/enable. Also, CDC jobs get deleted whenever we disable a VDB. The Post Start hook will add the CDC jobs again on enabling the VDB.

- e. Under VDB configuration, enable **Auto VDB Restart** to allow the Delphix Engine to automatically restart the VDB when it detects the target host reboot.
- f. To enable **Change Data Capture (CDC)**, select the Enable checkbox.
- g. Click **Next**.

 **CDC documentation**

For more information see:

- [About Change Data Capture](#)
- [Enable and Disable Change Data Capture](#)

10. Under **Policies**, select a **Snapshot Policy** for the VDB and click **Next**.
11. Under **Masking**, select **Mask this VDB** checkbox to mask your data during provisioning and then select one of the following masking options:
 - a. Select an existing masking job
 - b. Masking Job is not currently available for the selected data type, please mask using script(s) instead. If you select this option, you should define a Configure Clone script in the Hooks step to mask the dataset.
12. Under **Hooks**, specify any Hooks to be used during the provisioning process. For more information, see [Hooks for SQL Server](#).
13. Click **Next**.
14. The final summary tab will enable you to review your configurations.
15. Click **Submit**.

When provisioning starts, the VDB will appear in the **Datasets** panel. Select the VDB and navigate to the **Status** tab to see the progress of the job. When provisioning is complete, you can see more information on the **Configuration** tab.

 You can select a SQL Server instance that has a higher version than the source database and the VDB will be automatically upgraded. For more information about compatibility between different versions of SQL Server, see [SQL Server Support Matrix](#).

Provisioning by Snapshot or LogSync

When provisioning by snapshot, you can provision to the start of any particular snapshot, either by time or LSN.

 You can take a new snapshot of the dSource and provision from it by clicking the **Camera** icon.

Provisioning By Snapshot	Description
Provision by Time	You can provision to the start of any snapshot by selecting that snapshot card from the TimeFlow tab, or by selecting  and entering a value in the time entry fields. The values you enter will snap to the beginning of the nearest snapshot.

Provisioning By Snapshot	Description
Provision by LSN	You can use Provision by LSN control to open the LSN entry field. Here, you can type or paste in the LSN to which you want to provision. After entering a value, it will "snap" to the start of the closest appropriate snapshot.

Related Topics

- [SQL Server Support Matrix](#)
- [Hooks for SQL Server](#)

V2P with a SQL Server VDB

This topic describes how to perform the Virtual to Physical (V2P) process with a SQL Server virtual database (VDB).

Procedure

1. Log into the **Delphix Management** application.
2. Click **Manage** and select **Datasets**.
3. Select the dSource or VDB you want to export.
4. Select the snapshot of the dSource or VDB state you want to export.
5. If you want to export the state of the database from a specific point in time, select the **LogSync** icon and then select the point in time from which you want to create the export.



6. From the actions menu (...) select **Virtual to Physical**.
7. Select the target environment.
8. Enter the **Target Directory** for the export. The target environment should have been added to Delphix previously and should meet all target host requirements, see [Overview of Requirements for SQL Server Environments](#) for more information on user requirements for target environments. The target directory you enter here must exist in the target environment and the Delphix operating system user listed under the environment must have permission to write to it.
9. Select an option for **Run recovery after V2P**.
If you select **No**, you can use the scripts that are created in the target environment to manually recover the database at a later time. See [Manually Recovering a Database after V2P](#) for more information.
10. Click **Advanced** to customize the target directory layout or to enable Change Data Capture (CDC) on the exported database. See [Customizing Target Directory Structure for Database Export](#) for more information.
11. Click **Next**.
12. Select whether you want to have an email sent to you when the export process completes, and then click **Submit**.

Post-Requisites

If you selected **No** for **Run Recovery after V2P**, follow the instructions in [Manually Recovering a Database after V2P](#) to complete the V2P process.

Related Topics

- [Manually Recovering a Database after V2P](#)
- [Customizing Target Directory Structure for Database Export](#)
- [Overview of Requirements for SQL Server Environments](#)

Additional SQL Server VDB Topics

This section covers the following topics:

- [Renaming a SQL Server VDB via CLI](#)
- [Upgrading SQL Server VDBs](#)
- [Extended Properties for SQL Server VDBs](#)
- [File Permissions for SQL Server VDBs](#)
- [Deleting a SQL Server VDB](#)

Renaming a SQL Server VDB via CLI

This topic describes how to change the database name on the SQL Server Instance for VDB through the Delphix CLI.

Database name on SQL Server vs. VDB name on Delphix

The database name is what you would see the SQL Server instance on the Target in SQL Server Management Studio or sys.databases. It is also the database name in the **Configuration** tab in Delphix. The name of the VDB is an internal name within Delphix Engine and does not need to be the same as the database name on Target. This is found on the **Status** tab of Delphix.

Prerequisites

- The VDB must be running on the target environment.
- The SQL Server instance on the target environment where the VDB is located must be up and reachable.

Procedure

1. Select the **source** associated with the VDB and disable the VDB.

```
delphix> source "vexample"
delphix source 'vexample'> disable
delphix source 'vexample' disable *> commit
```

2. Select the source **config** associated with the source.

```
delphix source "vexample"> get config
vexample
delphix source "vexample"> /sourceconfig "vexample"
delphix sourceconfig "vexample">
```

3. Update the **databaseName** to the new name.

```
delphix sourceconfig "vexample"> update
delphix sourceconfig "vexample" update *> set databaseName=newDBName
delphix sourceconfig "vexample" update *> commit
delphix sourceconfig "vexample">cd
```

4. Enable the VDB.

```
delphix> source "vexample"
delphix source 'vexample'> enable
delphix source 'vexample' enable *> commit
```

Related Topics

- [Additional SQL Server VDB Topics](#)

Upgrading SQL Server VDBs

This topic describes how to upgrade a SQL Server VDB to a higher version of SQL Server instance.

Procedure for VDB In-Place Upgrade

1. Remove any VDB Refresh Policy assigned to the VDB.
2. Upgrade the target SQL Server instance.
3. Refresh the target environment.

Procedure to Upgrade a VDB to a New SQL Instance

1. Refresh all environments.
2. Login to the **Delphix Management** application.
3. Click **Manage**.
4. Select **Datasets**.
5. Select the **VDB** to be upgraded.
6. From the **Actions** menu (...) select **Disable**.
7. Click **Disable** to confirm.
8. From the **Actions** menu (...) select **upgrade**.
The **Upgrade Database** window will open.
9. Select the **SQL Server instance** to which you want the VDB to upgrade.
10. Click **OK**.
11. Enable the VDB. (See the *Enabling and Disabling SQL Server VDBs* section in [Provisioning and Managing Virtual Databases](#).)
12. Repeat steps 5 to 12 for each VDB you want to upgrade.

Related Topics

- [Additional SQL Server VDB Topics](#)

Extended Properties for SQL Server VDBs

Extended Properties and How to View Them

This topic describes extended properties on VDBs that you can use to track the origin of VDBs through SQL Server tools on target servers.

These are the extended properties:

Property	Description
<code>d\lpx_server_name</code>	Address of the Delphix Engine hosting the VDB
<code>d\lpx_server_uuid</code>	UUID of the Delphix Engine hosting the VDB
<code>d\lpx_source_id</code>	Internal reference of the VDB

You can find these properties in the following locations:

- For a VDB using the SQL Server Management Studio tool: under the Extended Properties page of the Properties window
- Using the `sp_d\lpx_vdbinfo` stored procedure.
 - To install and run this stored procedure, run the SQL code contained in `<Delphix Connector install path>\etc\sp_d\lpx_vdbinfo.sql`.

Related Topics

- [Additional SQL Server VDB Topics](#)

File Permissions for SQL Server VDBs

When provisioning a VDB, the Delphix Engine modifies the "access control lists" (ACLs) of database and log files to help prevent unintentional data loss through file deletion. Files could be deleted, for example, if there is an attempt to DROP a VDB directly through SQL Server management studio or other native SQL Server tools.

The Delphix Engine updates each database and log file ACL to include a deny-delete "access control entry" (ACE) for the user account running the SQL Server instance.

You can still drop VDBs directly through SQL Server tools. However, a warning message will be displayed, and the files will remain on the volume that the Delphix Engine exports. This file deletion prevention also applies to attempts to remove files from a database using the ALTER DATABASE .. REMOVE FILE command.

If a VDB is inadvertently dropped, you can reattach the database using SQL Server tools.

If you attempt to delete a database or log file and then try to add a file of the same name, this may fail because the original file was prevented from being deleted by the deny-delete ACE.

If you intend to delete the files from the volume that the Delphix Engine provides, you can change the ACLs on the files using the `icacls` command:

```
icacls <file> /remove <SQL Server instance account>:deny(D)
```

Accounts other than the SQL Server instance account will not be prevented from deleting the VDB database and log files.

Related Topics

- [Additional SQL Server VDB Topics](#)

Deleting a SQL Server VDB

Procedure

 Deleting a VDB is an unrecoverable operation. Proceed only if you want to permanently destroy the unique data that was created in the VDB.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **VDB** that you want to delete.
5. From the Actions menu (...) select **Delete**.
6. If stopping or starting the VDB requires particular credentials for the target environment other than those of the default environment user:
 - a. Check **Provide Privileged Credentials**.
 - b. Enter the **username** and **password**.
 - c. Click **Validate Credentials**.
7. Click **Delete** to confirm that you want to delete the VDB.


If the VDB was currently active, the Delphix Engine will shut it down, unmount all filesystems from the target environment, and finally delete the VDB itself.

Using Force Delete

Deleting unused or outdated objects should be a regular part of Delphix Engine administration. This is especially important to prevent low space errors, which can cause the Delphix Engine to stop. **The Delphix Engine holds a maximum of 400 objects.**

1. Log into the **Delphix Management** application.
2. Select **Resources > Storage Capacity**.
3. Next to the object you want to delete select the **Trashcan**.

- In the Delete dialog select **Force Delete**. Oracle users will have the option to provide additional credentials.



Delete Dataset Child VDB [X]

Are you sure you want to delete dataset "Child VDB"?

Force Delete


Provide privileged credentials

Cancel Delete

- Click Delete.


Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot. These items are displayed with a lock icon next to the name.



Deleting a VDB Associated with a Self-Service Container

As shown below if a VDB is associated with a Self-Service container, the delete option will not show.



Users will need to first delete the container and then delete the VDB. Refer to [Delphix Self-Service Data Container Activities](#).

Related Topics

- [Overview of SQL Server Virtual Databases](#)
- [Provisioning a SQL Server VDB](#)


CDC Support in SQL Server

Provisioning CDC Enabled Virtual Databases

Please see [Provisioning a SQL Server VDB](#).

Viewing or Updating CDC your Configuration

CDC configuration used while provisioning would be saved and used for subsequent Refresh and Rewind operations on the VDB. This configuration can be viewed/updated after the VDB has provisioned, in the ‘Source’ tab under the **Configuration**’section of the VDB.



Exporting CDC Enabled Physical Databases

1. In the Configuration page of the wizard, open the Advanced section and Select the **‘Enable’** option under ‘Change Data Capture (CDC)’

Virtual to Physical

Target Environment

Configuration

Summary

Simple

Open Database After Recovery

Hide advanced

Data Directory

Archive Directory

Temp Directory

External Directory

Script Directory

Change Data Capture (CDC) ⓘ

Enable

2. CDC capture and cleanup jobs have to be added and CDC metadata has to be upgraded (if exporting is done from a lower database version to SQL2016 and above), manually on exported databases.

General Notes/Troubleshooting

- Make sure the SQL Server instance on which the VDB is being provisioned or exported, supports CDC.
- Make sure 'SQL Server Agent' for the instance is running otherwise adding CDC capture and cleanup jobs will fail.

Related Topics

- [Provisioning a SQL Server VDB](#)

Hooks for SQL Server

Overview


This topic describes the use of hook operations with dSources created from SQL Server source databases and virtual databases (VDBs) that are created from SQL Server dSources or other VDBs.

Hooks are Windows Powershell code executed on:

1. The staging target host before or after the manual snapshot of a dSource.
2. The VDB target host before and after the provision, refresh, rewind, snapshot, start or stop of a VDB.

Powershell version

While creating a hook, a user can provide the PowerShell version (default version or version 2) in the field "Operation Type" and this PowerShell version will be used to execute the hook script. Here, the default version is the version of PowerShell installed on the target host.



Add a new Hook Operation ✕

Hook Point

Pre Sync Operations performed before a sync with external data.

Enter a Name

Enter a name

Operation Type

Run PowerShell Command (with default version)

 Run PowerShell Command (with version 2)

Enter command (or leave blank)

Cancel Create

Hooks can be specified in the wizard used during the creation of a VDB, or modified afterward by navigating to the Configuration > Hooks tab. Hooks can also be set using the Delphix command-line interface (CLI) or REST Web API.

Each hook operation represents a user-configurable action that the Delphix virtualization engine will execute. You can configure the custom hook code to fail if they encounter an unexpected error. The failure of a hook operation will cause the enclosing operation to fail.

The Windows environment user for the dSource or VDB runs the "Powershell" executable, which runs the specified PowerShell script on the Staging or VDB Target host. The Delphix Engine captures and logs all output of the script and displays it if a failure occurs.

The intent of hook operations are customization of the data contents or configuration of a dataset while it is being manipulated. Actions performed by hooks effectively become an integrated part of the sync operations of a dSource or the provision, refresh, rewind, snapshot, start, or stop actions for that VDB.

Hooks are mainly used for pre- and post-provisioning operations. For example, you can use hooks to:

- Back up test data before refresh and rewind
- Back up data after provisioning
- Reset configuration settings from production to non-production settings after provisioning
- Create logins for dev/qa users who do not have privileges on production databases
- Sync logins on the target that are cloned from the production database
- Back up configuration data from the database

For more information on Hook Operations, see [SQL Server Hook Operation Notes](#).

Hook Operation Templates

You can use operation templates to store commonly used operations, which allows you to avoid repeated code entry when an operation is applicable to more than a single hook, dSource, or virtual dataset. You can manage templates through the Delphix Management application.

You can also create templates from existing hooks by exporting the hooks in the Delphix Management application.

While creating a hook template, the user can provide the PowerShell version (default version or version 2) in the field "Type" and this PowerShell version will be used to execute the hook's script created from this template. Here, the default version is the version of PowerShell installed on the target host.

The screenshot shows a 'New Template' dialog box. It has a title bar with the text 'New Template' and a close button (X). Below the title bar, there are three input fields: 'Name', 'Description', and 'Content'. A dropdown menu is open over the 'Name' field, showing a list of options for 'Type': 'System Shell Command' (highlighted), 'Bash Shell Command', 'Expect Script', 'PowerShell Script (with default version)', and 'PowerShell Script (with version 2)'. At the bottom right of the dialog, there are two buttons: 'Cancel' and 'Create'.

The existing template's PowerShell version can be changed by using Delphix CLI only because UI currently does not support this feature.

Windows Environment Variables

When a hook is executed, Delphix will set specific Windows environment variables to provide context, such as the name of the current host, the name of the SQL Server instance and port, and the name of the database. For more information, see [SQL Server Hook Operation Notes](#).

Related Topics

- [SQL Server Hook Operation Notes](#)
- [Cookbook of Common Scripts for Hooks on SQL Server](#)
- [SQL Server Environments and Data Sources](#)
- [CLI Cookbook: Changing the PowerShell version associated with a hook template](#)

Python script to migrate hooks from PowerShell version 2 to host's Default PowerShell version.

Overview

With 6.0.3.0, the Delphix Engine will use the default PowerShell version installed on the host (hereinafter referred as default PowerShell) to perform all its operations, also the new hooks and hook templates can be created using the default PowerShell.

The existing hooks and hook templates on the engine will continue functioning using PowerShell version 2, and there will be an option to migrate them to use default PowerShell.

The article intends to introduce a Python script that can be used to migrate all the hooks and hook templates on the engine to default PowerShell. The motivation behind writing the script is to save the manual effort required in doing the migration via the UI or CLI.

[Link to download the script.](#)

Requirements for running the script

The requirements for running the script are as follows, they are the same as running any Python script in general.

1. The machine where the script is run should have Python installed.
 - a. Relevant link: <https://www.python.org/downloads/>
 - b. The script is supported for both Python 2 and Python 3 release.
2. The machine where the script is run should have the 'delphixpy' Python package installed.
 - a. Relevant link: <https://pypi.org/project/delphixpy/>
 - b. [python.org](#) documentation on installing packages and creating Python environments: [Installing packages using pip and virtual environments.](#)
 - c. If the package is already installed, it should be upgraded to the latest version, the minimum *delphixpy* version required to run the script is 1.11.3.0
3. The Delphix Engine should be accessible from the machine where the script is run since the script makes API requests to the Delphix Engine to perform the migration,
 - a. The easiest way to verify the same is to use the **ping** command.

Functionalities of the script

1. The script can migrate the hooks and hook templates in the Delphix Engine to run with default PowerShell.
2. The script can also migrate the hooks and hook templates back to PowerShell version 2.
3. It is possible to migrate only the hooks or only the hook templates.
4. If **INSTALLEDPOWERSHELL** feature flag is enabled on the Delphix Engine, the script will disable the same.

The script parameters and usage

Script parameter	Description	Type	Possible values	Default value
--help	Displays the description and usage details of the parameters for the script on the terminal	Optional	Not applicable	Not applicable
--engine-addr	The Delphix Engine host address.	Required	Not applicable	Not applicable

Script parameter	Description	Type	Possible values	Default value
--sys-admin-usr	The username for System Administrator user to log into the Delphix Engine.	Required	Not applicable	Not applicable
--sys-admin-pwd	The password for System Administrator user to log into the Delphix Engine.	Required	Not applicable	Not applicable
--admin-usr	The username for Delphix Administrator user to log into the Delphix Engine.	Required	Not applicable	Not applicable
--admin-pwd	The password for Delphix Administrator user to log into the Delphix Engine.	Required	Not applicable	Not applicable
--hook-ps-version	Migrates all the hooks to PowerShell Version two.	Optional	default, ps2	default
--hook-templ-ps-version	Migrates all the hook templates to PowerShell Version two.	Optional	default, ps2	default
--migrate-only	Migrates only the hook or only the hook templates. Input "hooks" for migrating only the hooks, whereas input "templates" for migrating only the hook templates	Optional	hooks, templates, hooks-and-templates	hooks-and-templates
--debug	In case there is an error executing the script, prints the Python stack trace required for debugging.	Optional	Not applicable	Not applicable

Command examples

1. Displaying the description and usage details of the parameters for the script on the terminal

```
python <path to the script> --help
```

2. An example command for migrating all the hook and hook templates to default PowerShell version:

```
python <path to the script> --engine-addr engine.delphix.com --sys-admin-usr sysadmin --sys-admin-pwd sysadmin --admin-usr admin --admin-pwd delphix
```

3. An example command for migrating all the hook and hook templates to PowerShell version 2:

```
python <path to the script> --engine-addr engine.delphix.com --sys-admin-usr
sysadmin --sys-admin-pwd sysadmin --admin-usr admin --admin-pwd delphix --hook-
ps-version ps2 --hook-templ-ps-version ps2
```

Please note that this will migrate all the hooks and hook templates in the Delphix Engine to PowerShell version two, the script does not have the functionality to perform the migration for specific hooks or hook templates.

4. An example command for migrating only the hook templates:

```
python <path to the script> --engine-addr engine.delphix.com --sys-admin-usr
sysadmin --sys-admin-pwd sysadmin --admin-usr admin --admin-pwd delphix --
migrate-only templates
```

5. An example command for running the script in debug mode:

```
python <path to the script> --engine-addr engine.delphix.com --sys-admin-usr
sysadmin --sys-admin-pwd sysadmin --admin-usr admin --admin-pwd delphix --debug
```

Common Errors while running the script:

1. If the Delphix engine is not accessible from the machine where the script is run, the connectivity test run by the script prior to running the migration will fail, and the following error will be thrown:

```
Error occurred while connecting to the engine via the given Delphix System
Administrator user:
[Errno 8] nodename nor servname provided, or not known
```

```
Error occurred while connecting to the engine via the given Delphix
Administrator user:
[Errno 8] nodename nor servname provided, or not known
```

2. The connectivity test run by the script prior to running the migration will fail if the value of the required parameters provided is not correct, and the following error will be thrown:

```
Error occurred while connecting to the engine via the given Delphix
Administrator user:
HTTP status was 401 when doing POST '{"username": "admin", "password": "blah",
"type": "LoginRequest", "target": "DOMAIN"}' to '/resources/json/delphix/
login': {"type":"ErrorResult","status":"ERROR","error":
{"type":"APIError","details":"Invalid username or password.,"action":"Try with
a different set of
credentials.,"id":"exception.webservices.login.failed","commandOutput":null,"d
iagnoses":[]}}
```

3. If for some reason, the script execution stops in between before migrating all the hooks and hook templates, it's perfectly fine to run the script again; each time the script is run, the hooks and hook templates will be migrated to the specified PowerShell version.

SQL Server Hook Operation Notes

SQL Server Clusters

When linking from, or provisioning to cluster environments, hook operations will not run once on each node in the cluster. Instead, the Delphix Engine always runs all hooks on the instance primary node.

RunPowershell Operation

The RunPowershell operation executes a PowerShell script on a Windows environment. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output of the script. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunPowershell Operation

You can input the full command contents into the RunPowershell operation.

```
$removedir = $Env:DIRECTORY_TO_REMOVE

if ((Test-Path $removedir) -And (Get-Item $removedir) -is [System.IO.DirectoryInfo])
{
    Remove-Item -Recurse -Force $removedir
} else {
    exit 1
}
exit 0
```

SQL Server Environment Variables

Operations that run user-provided scripts have access to environment variables. For operations associated with specific dSources or virtual databases (VDBs), the Delphix Engine will always set environment variables so that the user-provided operations can use them to access the dSource or VDB.

dSource Environment Variables

Environment Variables	Description
<code>SOURCE_INSTANCE_HOST</code>	The hostname of linked instance for the dSource
<code>SOURCE_INSTANCE_PORT</code>	Port of linked instance for the dSource
<code>SOURCE_INSTANCE_NAME</code>	Name of linked instance for the dSource
<code>SOURCE_DATABASE_NAME</code>	Name of database linked for the dSource

Staging Variables

We have the following environment variables applicable to Staging Push dSources.

Environment Variables	Description
STAGING_INSTANCE_HOST	The hostname of the staging instance
STAGING_INSTANCE_PORT	Port number of the staging instance
STAGING_INSTANCE_NAME	Name of the staging instance
STAGING_DATABASE_NAME	Name of the staging database
STAGING_MOUNT_BASE	Mount path for the staging push dSource
STAGING_DATA_DB_FILE_PATH	Filepath of the staging database

VDB Environment Variables

Environment Variables	Description
VDB_INSTANCE_HOST	The hostname of linked instance for the VDB
VDB_INSTANCE_PORT	Port of linked instance for the VDB
VDB_INSTANCE_NAME	Name of linked instance for the VDB
VDB_DATABASE_NAME	Name of database linked for the VDB

Error handling for SQL Server PowerShell Scripts

If a pre-script or post-script encounters an unrecoverable error during execution, the Delphix Engine expects the script to return with a non-zero exit code or the error will not be detected. The **Powershell -File** prefix and **exit \$LASTEXITCODE** suffix are required to pass the script's exit code up to the layer calling the script.

 Delphix does not perform error checking on PowerShell hook scripts. The script should perform error checking and logging, and return a non-zero exit code to indicate the script's failure. Failure to return a non-zero exit code when appropriate means that Delphix will think the hook script succeeded and mark

the VDB provision/refresh/rewind job as a success, when it should be seen as a failure. This is especially important when masking data is part of the hook – the VDB should not be released to users when the hook failed to mask data.

PowerShell gives you a few ways to handle errors in your scripts:

- Set `$ ErrorActionPreference` . This only applies to PowerShell Cmdlets. For scripts or other executables such as `sqlcmd` , PowerShell will return with exit code 0 even if there is an error, regardless of the value of `$ErrorActionPreference` . The allowable values for `$ErrorActionPreference` are:
 - `Continue` (default) – Continue even if there is an error
 - `SilentlyContinue` – Acts like Continue with the exception that errors are not displayed
 - `Inquire` – Prompts the user in case of error
 - `Stop` : Stops execution after the first error
- Use exception handling by using traps and try/catch blocks to detect errors and return with non-zero exit codes
- Use custom error handling that can be invoked after launching each command in the script to correctly detect errors. The following example shows how you can use the function `verifySuccess` to detect whether the previous command failed, and if it did print, print an error message and return with an exit code of **1**.

```
function die {
    Write-Error "Error: $($args[0])"
    exit 1
}

function verifySuccess {
    if (!$?) {
        die "$($args[0])"
    }
}

Write-Output "I'd rather be in Hawaii"
verifySuccess "WRITE_OUTPUT_FAILED"

& C:\Program Files\Delphix\scripts\myscript.ps1
verifySuccess "MY_SCRIPT_FAILED"
```

Related Topics

- [SQL Server Environments and Data Sources](#)
- [Cookbook of Common Scripts for Hooks on SQL Server](#)

Using Pre- and Post-Scripts with SQL Server dSources

Overview

This topic describes the use of pre- and post-scripts with dSources that are created from SQL Server source databases.

Pre-scripts and post-scripts are Windows PowerShell code executed on the Staging Target host before and after a SnapSync of a dSource. You can specify pre- and post-scripts in the wizard for creating a dSource, or you can modify them afterward by navigating to the Configuration > Standard tab. You can also set pre- and post-scripts using the Delphix command-line interface (CLI) or REST Web API.

The Delphix Engine executes a pre-script on the Staging Target host prior to the SnapSync of a dSource. If it is an initial snapshot, or manual snapshot by selecting the snapshot button on the GUI, the pre/post scripts do not get executed for dSources. Since the dSource resides within a restoring database in the SQL instance on the Staging Target host, the script can perform queries on the instance if a database account is available.

Hooks do allow for a pre/post snapshot hook, but the pre/post scripts do not

A post-script is executed after the SnapSync on a dSource completes. If the post-script fails, the provision, refresh, or rewind operation will also fail with an error message, and a fault will be created on the dSource.

 Pre/Post hooks allow for pre/post snapshot hooks for which there is no equivalent in the old pre/post scripts. Pre/Post scripts are run during validated sync, for which there is no equivalent hook operation today. Pre- and post-scripts are supported for backward compatibility with older versions of the Delphix Engine.

Associating Scripts with a dSource

Pre- and Post-scripts can be associated with a dSource in one of two ways:

During the dSource linking process

1. Login to the **Delphix Management** application.
2. In the top menu bar, click **Manage**.
3. Select **Datasets**.
4. Select the **Plus** icon and then select **Add dSource**.
5. In the **Hooks** tab of the **Add dSource** wizard, select the Plus icon to add new **Pre Script** and **PostScript** hooks.
6. Enter the calling syntax of the Windows PowerShell script into either or both of the appropriate fields.
 - a. The calling environment is used during linking, as shown in the **Environment Details** panel of the Delphix environment (**Manage > Environments**)
 - b. Four (4) environment variables will be populated with the name of the Delphix dSource, the SQL Server instance name and port, and the SQL Server database
 - c. name.

After linking, using the Configuration tab of the Datasets details page

1. Login to the **Delphix Management** application.
2. In the top menu bar, click **Manage**.
3. Select **Datasets** to display the SQL Server dSources and VDBs.
4. Select a SQL Server **dSource** from the listed **Datasets** in the left-hand navigation bar.
5. Click the **Configuration** panel, select the Hooks sub-tab.

6. By selecting a Plus icon you can create hooks from a template or create a new hook.

7. Enter the calling syntax of the Windows Powershell script into either or both of the appropriate fields
- The calling environment is that of the **primary Environment User account**, as shown in the **Environment Details** panel of the Delphix environment (**Manage > Environments**)
 - Four (4) environment variables will be populated with the name of the VDB, the SQL Server instance name and port, and the SQL Server database name.
 - Select the **Create** to accept the change.

Execution Context for SQL Server Scripts

For dSources, pre- and post-scripts are executed in the context of the staging host user that was provided when linking.

Related Topics

- [Hooks for SQL Server](#)

Using Pre- and Post-Scripts with SQL Server VDBs

Overview

This topic describes the use of pre- and post-scripts with virtual databases (VDBs) that are created from SQL Server dSources.

Pre-scripts and post-scripts are Windows PowerShell code executed on the VDB target host before and after the provision, refresh, or rewind of a VDB. You can specify pre- and post-scripts in the wizard for creating a VDB, or you can modify them afterward by navigating to the Configuration > Standard tab. You can also set pre- and post-scripts using the Delphix command-line interface (CLI) or REST Web API.

The intent of these scripts is a customization of the data contents or configuration of a VDB while it is being manipulated. Actions performed by pre-scripts and post-scripts effectively become an integrated part of the provision, refresh, or rewind actions for that VDB.

The pre-script executes during the initial provision of a VDB. During refresh and rewind, the PowerShell script referenced in a pre-script is executed after the VDB has been stopped and unmounted, but before the new VDB is mounted. If the pre-script fails, the refresh or rewind operation will also fail with an error message.

During provision, refresh, and rewind, the PowerShell script referenced in a post-script is executed after the Delphix engine has mounted and started the VDB. If the post-script fails, the provision, refresh, or rewind operation will also fail with an error message, and a fault will be created on the VDB.

You can use a pre-script to capture configuration file settings, but not the contents of the soon-to-be recreated VDB; a pre-script executes too late to access the VDB which has already been shut down and unmounted. This makes pre-script functionality much less useful than hook operations like Pre-Refresh.

You can use a post-script to run data transformation operations on newly-provisioned, newly-refreshed, or newly-rewound VDBs. These operations include data masking and setting non-production account passwords in place of cloned production passwords.

 Pre- and post-scripts are an older customization mechanism for SQL Server virtual databases. They have been replaced by hook operations, which have been the standard customization mechanism on all other data platforms. Pre- and post-scripts are supported for backward compatibility with older versions of the Delphix Engine. Delphix encourages everyone to use Hooks for customizing SQL Server VDBs for future implementations, if possible.

Associating Scripts with a VDB

Pre- and Post-scripts can be associated with a VDB in one of two ways:


During the VDB provisioning process

1. Login to the **Delphix Management** application/
2. In the top menu bar, click **Manage**.
3. Select **Datasets** to display the SQL Server dSources and VDBs.
4. From the listed **Datasets** in the left-hand navigation bar, select a SQL Server **dSource** or **VDB**.
5. Click the **TimeFlow**.
6. Click **Provision**.
7. In the first **Target Environment** step of the **Provision VDB** wizard, there are fields for **Pre Script** and **Post Script**.
8. Enter the calling syntax of the Windows PowerShell script into either or both of the appropriate fields.

- a. The calling environment is that of the **primary Environment User account**, as shown in the **Environment Details** panel of the Delphix environment (**Manage > Environments**)
- b. Four (4) environment variables will be populated with the name of the VDB, the SQL Server instance name and port, and the SQL Server database name.

After provisioning, using the Configuration tab of the Datasets details page

1. In the **Datasets** panel, click the virtual dataset.
2. Click the **Configuration** tab.
3. Within the **Configuration** tab, click the **Hooks** tab.
4. Select the **hook** to edit.
 - a. Click the **Plus** icon to add a new operation.
 - b. Select the **type of operation** or click  to load a hook operation template.
 - c. Click the **text area** and edit the contents of the operation.
 - d. To remove an operation from the list, click the **Trash** icon on the operation.
 - e. When you have set all hook operations, click the **checkmark** to save the changes. The current operations at this hook will be displayed. To edit this list of operations, click the **Pencil** icon in the top right-hand corner of the tab.



Execution Context for SQL Server Scripts

For VDBs, pre- and post-scripts are executed in the context of the environment user that was selected during the VDB provision and not the primary environment user. The primary environment user can change over time, all VDB operations are done using the user that was initially selected.

Error handling for SQL Server PowerShell Scripts

If a pre-script or post-script encounters an unrecoverable error during execution, the Delphix Engine expects the script to return with a non-zero exit code. Otherwise, the error will not be detected.

PowerShell gives you a few ways to handle errors in your scripts:

- Set `$ ErrorActionPreference` . This only applies to PowerShell Cmdlets. For scripts or other executables such as `sqlcmd` , PowerShell will return with exit code 0 even if there is an error, regardless of the value of `$ErrorActionPreference` . The allowable values for `$ErrorActionPreference` are:
 - `Continue` (default) – Continue even if there is an error
 - `SilentlyContinue` – Acts like Continue with the exception that errors are not displayed
 - `Inquire` – Prompts the user in case of error
 - `Stop` : Stops execution after the first error

- Use exception handling by using traps and try/catch blocks to detect errors and return with non-zero exit codes
- Use custom error handling that can be invoked after launching each command in the script to correctly detect errors. The following example shows how you can use the function `verifySuccess` to detect whether the previous command failed, and if it did print, print an error message and return with an exit code of **1**.

```
function die {
    Write-Error "Error: $($args[0])"
    exit 1
}

function verifySuccess {
    if (!$?) {
        die "$($args[0])"
    }
}

Write-Output "I'd rather be in Hawaii"
verifySuccess "WRITE_OUTPUT_FAILED"

& C:\Program Files\Delphix\scripts\myscript.ps1
verifySuccess "MY_SCRIPT_FAILED"
```

Related Topics

- [Hooks for SQL Server](#)

Using HostChecker to Validate Target Database Servers

Prerequisites

Make sure that your source and target environments meet the requirements.

Procedure to Validate Target Environments

1. Verify with your System Administrator that the Delphix Connector has been installed in all Target environments.
2. Login to the Windows target host using the **Windows user account** that the System Administrator configured as a Delphix target user.
3. Open **Windows Powershell** using the **Run as Administrator** option.
4. Execute the **host checker script** by running:
`<Delphix Connector installation folder>\etc\dlpx-host-checker.ps1`
5. Select a path where a report file will be saved, such as **C:\temp\delphix-host-checker-report.txt**.
6. Select the **default** option of **Target Host**.
7. Read the output of the checks.
8. The error or warning messages will explain any possible problems and how to address them. Resolve the issues that the HostChecker describes. Do not be surprised or undo your work if more errors appear the next time you run HostChecker; the error you just fixed may have been masking other problems.
9. Repeat steps 4–7 until all the checks return no errors or warnings.

Tests Run

Test	SQL Server Source	SQL Server Target	Description
Check Powershell Version	X	X	Verifies that Powershell 2.0 or greater is installed
Check OS User Privileges	X	X	For target hosts, verifies that the operating system (OS) user has administrative rights. For source hosts, verifies that the OS user can successfully perform remote registry access from the target host to the source host.
Check host settings	X	X	Verifies that the Delphix Engine can discover host environment details from the Windows registry.
Check SQL Server instance discovery	X	X	Verifies that the Delphix Engine can discover SQL Server instances.

Test	SQL Server Source	SQL Server Target	Description
Check SQL Server instance login permission	X	X	For target hosts, verifies that the Windows OS user can be used to log in to the SQL Server instances. For source hosts, verifies that the supplied SQL Server login credentials can be used to log in to the SQL Server instances.
Check database discovery	X	X	Verifies that the Delphix Engine can discover SQL Server databases.

Additional Options

Run the following to view additional HostChecker options:

```
d\px-host-checker.ps1 -?
```

Related Topics

- [SQL Server Support and Requirements](#)

SAP ASE Environments and Data Sources


This section covers the following topics:

- [Delphix Architecture with SAP ASE](#)
- [Overview of ASE Database Encryption](#)
- [Quick Start Guide for SAP ASE](#)
- [SAP ASE Support and Requirements](#)
- [Managing SAP ASE Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Provisioning and Managing VDBs from SAP ASE](#)
- [Backup Server Best Practices](#)
- [SAP ASE Hook Operations](#)
- [Support for Dump History File](#)

Delphix Architecture with SAP ASE

This topic describes the high-level process for adding SAP ASE-supported environments, linking SAP ASE databases to the Delphix Engine, and provisioning virtual databases.

Diagram of underlying Linking Architecture to support ingestion workflows between SAP ASE-supported Environments and the Delphix Engine.




Linking Architecture between SAP ASE and Delphix Engine

The diagram featured above shows a common architecture for SAP ASE-based systems. The diagram shows the infrastructure used to ingest data. (From the above diagram) The workflow starts with a Source Host (Left, which in this example is a Production Database Server), that provides critical data in the form of dump files (used for SnapSync) and transaction logs (used for LogSync).

The Delphix Engine (Bottom), continuously monitors the source database to determine when new dumps are available. When a dump is available, the Delphix Engine will contact the SAP ASE Staging host (Right) via Delphix Session Protocol (DSP). The Staging host will read the dumps from the Source Host and recover them through a Staging Database that is automatically set up on Delphix NFS storage. Once recovery is complete, the backup data

is incorporated into the Delphix dSource as a new snapshot card and is available for use to provision a new virtual or physical database.

Diagram of underlying Provisioning Architecture to support VDB Provisioning workflows between the Delphix Engine and SAP ASE-Supported Target Environments.



⚠ If no remote load location (aka Source Host) is configured, dump files must be made available on the staging host via NFS or some other copy mechanism. If no remote load location is configured, the Delphix Engine searches for the files on the staging server. If they are not found, the Delphix Engine cannot load them

SAP ASE Provisioning Architecture

This diagram is an extension of the previous diagram with the new content showing how Delphix provisions Virtual Databases (VDBs) to a Target Environment. The Delphix Engine (bottom) creates a set of virtual files from a snapshot that becomes the VDB. These files are presented to the Target host (right) via NFS. Delphix uses DSP to communicate with Target and initiates the creation of a new database. Once complete, the VDB is brought online and made available for use.

Environment Setup

Target and Staging Environments for SAP ASE

For Sap ASE-based Delphix Architectures, the Staging Host plays a critical role:

1. Provide the “staging” point in which the Delphix Engine coordinates data ingestion to a dSource by restoring backups to a staging database and creating a dSource from that ingested data.
2. Tracks changes on disk by running validated sync, a process that identifies when new backups are available and initiates the ingestion process if new backups and/or transaction log backups are found.
3. Minimizes touch to production by providing an intermediate host between source databases and Delphix.

For SAP ASE-based Delphix Architectures, the Target Host has two potential roles:

1. Host a target environment for provisioning virtual databases (VDBs).

 Staging and Target hosts can be run on a target server, although a dedicated staging server is recommended for optimal performance.

SAP ASE dSources are representations of a Source Database replica on the staging database that runs on a target or staging host. There is no requirement for additional local storage with either host option, as the storage is mounted over NFS from the Delphix Engine. For a deeper, technical discussion, please see the Technical Deep Dive section below.

At Delphix, we refer to the creation and maintenance of this staging database on the staging host as "validated sync," because it prepares the dSource data on the Delphix Engine for provisioning VDBs later on. After the Delphix Engine creates the staging database, it continuously monitors the source database for new transaction log/Full backups (if “truncate log on checkpoint” is active, Delphix will need Full Backups for Validated Sync). When it detects a new transaction log backup, it restores that backup to the staging database. The result is a TimeFlow with consistent points from which you can provision a VDB, and a faster provisioning process, because there is no need for any database recovery during provisioning. If Log Sync is enabled, you may provision a VDB to a point in time, in between snapshots.

When you later provision a VDB, you can specify any environment as a target, including the environment that contains the staging database. However, for best performance, we recommend that you choose a different target environment. The only requirement for the target is:

- It must have an operating system that is compatible with the one running on the validated host.

TARGET HOSTS for ASE

Container for VDBs

This topic describes the basic concepts involved with provisioning VDBs from SAP ASE dSources or even other SAP ASE VDBs.

A dSource is a virtualized representation of a physical or logical source database. As a virtual representation, it cannot be accessed or manipulated using database tools. Instead, you must create a virtual database (VDB) from a dSource snapshot. A VDB is an independent, writable copy of a dSource snapshot. You can also create VDBs from other VDBs. Once you have provisioned a VDB to a target environment, you can also implement snapshot and retention policies for the VDB, which will determine how frequently Delphix Engine will take a database snapshot and how long the snapshots will be retained for recovery and provisioning purposes.

When provisioning a VDB, Delphix creates the database with the default SAP ASE database options. If the database options have been altered on the source database and you wish for the VDB to reflect these same options, they would need to be altered via a Post-Script or by [Hook Scripts for Automation and Customization](#).

For an overview of the high-level components involved in provisioning an SAP ASE VDB, see [Overview of Provisioning SAP ASE Virtual Databases](#).

Validated Sync and LogSync

To run validated sync, a staging environment must be specified to host a staging database for the validated sync process. In this process, the Delphix Engine continuously monitors the source database for new full and transaction log backups if the source database is using a simple recovery model, and only transaction log backups if using a full recovery model. When it detects a new backup, it restores that backup to the staging database with the storage residing in Delphix. The result is a Timeflow with consistent points from which you can provision a VDB, also known as snapshots.

 If the “truncate log on checkpoint” is set, Delphix will only apply full backups.

Snapshots accumulate over time. To view a snapshot:

1. From the Datasets panel, click the group containing the dSource.
2. Select dSource.
3. Click the Timeflow tab.

Each snapshot is displayed and includes information about the source database, operating system, and time stamp. You can scroll through these cards to select the one you want, or you can enter a date and time to search for a specific snapshot.

Once you have provisioned a VDB, you can also take snapshots of it. As with the dSource snapshots, you can find these when you select the VDB in the Datasets panel. You can then provision additional VDBs from these VDB snapshots.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot.

Related Topics

- [SAP ASE Environments and Data Sources](#)

Overview of ASE Database Encryption

Beginning with 16.0, SAP ASE supports the Database Encryption feature. The SAP ASE Database Encryption feature encrypts the data at rest, without changing the applications. This encryption can be done on entire databases or only on columns. This ensures that the authorized users access the data and thus prevents the misuse of the data against theft and security breaches.

Data is encrypted with the help of encryption keys. These encryption keys are stored in the database in an encrypted form. You can encrypt an encryption key using a key encryption key (KEK).

In the SAP ASE Database Encryption, column and database encryption uses a symmetric encryption algorithm, which means that the same key is used for encryption and decryption. SAP ASE tracks the key that encrypts the data.

Starting 6.0.8.0, Delphix Engine will support the SAP ASE encrypted databases.

For more information on SAP ASE Database Encryption, see the [SAP ASE Encryption Documentation](#).

Delphix Implementation of Database Encryption

This topic describes various configurations to support encrypted databases with Delphix. Follow the mandatory steps below on the ASE instance that hosts the staging databases and virtual databases.

1. If the source database is not encrypted already.
 - a. Install the license option ASE_ENCRYPTION.
 - b. Create a master key that will serve as the KEK.

Command

```
> create encryption key master with passwd "sybase"
```

- c. If the database is not encrypted already, create the database encryption key and use it.

Commands

```
> create encryption key <encryption-key-name> for database encryption
> sp_configure "number of worker processes", 2
> alter database <database-name> encrypt with <encryption-key-name>
```

- d. Export the master key and the encryption key to a location that is shared among source, staging, and target hosts. The command-line version of the `ddlgen` tool is located at `$SYBASE/$SYBASE_ASE/bin`. You need to find out this location for your instance if it is different.
 - i. `cd $SYBASE/$SYBASE_ASE/bin`
 - ii. `ddlgen -Usa -Psybase -SASE160_SRC -TEK -N master.dbo.master -XOD -O<shared-path>/master_ddl.sql`
 - iii. `ddlgen -Usa -Psybase -SASE160_SRC -TEK -N master.dbo.<key_name> -XOD -O<shared-path>/<key_name>_ddl.sql`
- e. Enable encryption in SAP ASE by executing the below command on the staging/target instance.

Command

```
> sp_configure 'enable encrypted columns', 1
```

2. Import the keys on the staging and target instances by running the below commands from the directory where the SQL files are present (the shared location between instances) or mention the entire path of the files to be imported.
 - a. `isql -Usa -Psybase -SASE160_TGT -w 220 -i master_key.sql`
 - b. `isql -Usa -Psybase -SASE160_TGT -w 220 -i <key_name>.sql`
3. Set the encryption password by executing the below command on the staging/target instance.

Command

```
> set encryption passwd "sybase" for key master
```

4. Setup for automatic master key access. Refer [create the master key start-up file](#). In order to avoid issues on the master key password after the reboot of the ASE instance, a master key startup file needs to be created by running the following steps on the staging and the target instance.

a. **Command**

```
> sp_configure 'automatic master key access',1
```

b. **Command**

```
> alter encryption key master with passwd 'sybase' add encryption for
automatic_startup
```

c. **Command**

```
> sp_encryption mkey_startup_file,default_location,sync_with_mem
```

- d. Verify if the master key startup file has been successfully created on the instance.

Command

```
> sp_encryption mkey_startup_file
```

- e. Reboot the ASE instance to get the master key startup file in effect.

 If you perform a reboot or plan to perform a reboot on the source host, then you would need to repeat step 4 on the source host as well. By doing so, you don't need to set the master key password again after reboot.


Quick Start Guide for SAP ASE

This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a focused overview of working with SAP ASE database objects in the Delphix Dynamic Data Platform. It does not cover advanced configuration options. It does not cover advanced configuration options or best practices for performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix Engine functionality. It assumes you will use the VMware Hypervisor.

Overview

In this guide, we will walk through deploying a Delphix Engine, starting with configuring SAP ASE Source and Target environments. We will then create a dSource, and provision a VDB. A detailed list can be referred to in the [Network and Connectivity Requirements for SAP ASE Environments](#) section of the manual.

For purposes of the QuickStart, you can ignore any references to Replication or Masking, such as the engines shown in the diagram below.




Delphix stays in sync with source databases by monitoring the ASE backup server log. When it sees a new database dump or transaction log has been created, it attempts to load it into the ASE staging instance.

Color	Supported?
✓	Yes
✗	No

Validated sync Mode	Snapshot for FULL BACKUP	Snapshot for TLOGS	Point-in-Time Restore
truncate log on chkpt = true	✓	✗	✗
truncate log on chkpt = false	✗	✓	✗
truncate log on chkpt = true + Log Sync	✓	✗	✗
truncate log on chkpt = false + LogSync	✗	✓	✓

Before proceeding with setting up environments, decide where the databases will be located relative to the ASE staging instance. In the following diagram, we have two ASE hosts (production and staging).



If the database dumps (or transaction logs) are available locally to the staging Backup Server, Delphix considers them “local”. Whereas if they are available to a Backup Server on production, they are “remote”. If the files will be local to the staging server, there are many options to get them from the production host to the staging host including but not limited to:

- Sharing the directory on the production host over NFS to the staging host.
- Using Network Attached Storage (NAS) to replicate the files.
- Using scp to copy the files from one host to the other.

If the files are only available on the production host, Delphix will login to the ASE staging instance and issue the “LOAD DATABASE” command using the remote server syntax (for example LOAD DATABASE pubs2 FROM "/dumps/pubs2.full.9_5_16" AT Prod_BS).

Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online.
This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

```

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be
configured for networking in your
environment. From here, you can configure the primary interface, DNS, hostname,
and default route. When DHCP is
configured, all other properties are derived from DHCP settings.

To see the current settings, run "get." To change a property, run "set
<property>=<value>." To commit your changes,
run "commit." To exit this setup and return to the standard CLI, run "discard."

    defaultRoute    IP address of the gateway for the default route -- for
example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
the primary interface. Setting this value
                    to "true" will cause all other properties (address,
hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
"1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
-- for example, "1.2.3.4/22"

Current settings:

defaultRoute: 192.168.1.1
dhcp: false
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24

```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

8. Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`)

9. Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

10. Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the
browser at:
```

```
http://<address>
```

Type `"exit"` to disconnect, or any other commands to **continue** using the CLI.


11. Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
12. Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Once you set up the network access for Delphix Engine, navigate to the Delphix Engine URL in your browser for server setup.

The welcome screen below will appear for you to begin your Delphix Engine setup.



The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at `http://<Delphix Engine>/login/index.html#serverSetup`.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.

- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

⚠ The default domain user created on Delphix Engines from 5.3.1 is known as **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

System Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications.

Choose your option to set up system time in this section. For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click **Next** to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.

4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

⚠ Although your Delphix Engine will work without registration, we strongly recommend that you register each Delphix Engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

Requirements for SAP ASE Hosts and Databases

In order to begin using SAP ASE environments with Delphix, you will need to configure the source and target hosts with the requirements described on this page.

SAP ASE Source Host Requirements

There must be an operating system user, such as `delphix_os`, that meets the following requirements:

- The `$SYBASE` environment variable is defined for non-interactive shells (such as via the `.bashrc` configuration file).
 1. Set the `PermitUserEnvironment` configuration parameter to "yes" in the `sshd_config` file
 2. Add the variable to the user's `.ssh/environment` file
 3. Restart the SSH daemon
- To test this requirement:

```
ssh delphix_user@ase_hostname env | grep SYBASE
```

- Can login to the source host via SSH (TCP port 22)
- Delphix requires superuser permission to run **pargs** in order to discover Solaris ASE instances. For more information, see [Sudo Privilege Requirements for SAP ASE Environments](#).
- Designating the Delphix operating system user's primary group to be the same as the ASE instance's means the file system permissions can be more restrictive and is a better security practice than granting world read access to the toolkit or the backup files. If the target host is used to host the staging databases, consider the following:
 - If you don't add the Delphix operating system user to the ASE instance owner's group, greater permissions will need to be **granted to the backup files to ensure read access to the dumps and/or transaction logs**. Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment).

- If you sync Delphix with a dSource by asking ASE to create a new backup, the ASE instance owner will need the **write permission to the toolkit** (or the mount point if you use the CLI to specify a directory other than the toolkit). Delphix will issue the "DUMP DATABASE" command to write to the staging database's "temp" directory which is mounted on the staging host.
- Has **write permission for the mount-point directory** (by default the toolkit directory but can be a separate mount point specified in the command line interface).
- There must be a directory on the source host where you can install the Delphix platform toolkit, for example: /var/opt/delphix/Toolkit
- The delphix_os user must own the directory
- The directory must have permissions 0770, for example, -rwxrwx---. However, you can also use more permissive settings.
- The directory should have 256MB of available storage.

Source Database Requirements

When adding a source ASE environment to Delphix, you may use a single login to discover the ASE instances and link the source databases OR you may use a single login to discover all of the ASE instances and separate logins to link each dSource.

- Delphix uses a single database user for the discovery of all ASE instances and their databases for each environment added to Delphix.
 - The discovery database user (delphix_disc for example) must have SELECT privileges on the following tables for each ASE instance on the source host:
 - **sysdatabases**
 - **syservers**
 - **syslisteners**
 - **sysconfigures**
 - **syscurconfigs**
- Another user must be specified when linking each dSource (delphix_link for example) that has SELECT privileges on the above tables.
- If you will select New Full Backup when linking, this user must also have privileges to take a new full database dump of the source database. For more information about linking options, see [Linking an SAP ASE Data Source](#).
- The link database user can be different for each instance and database on the source host.

 If the source database is resized and trunc log on chkpt is disabled, take a transaction log dump immediately after the resize operation completes. If trunc log on chkpt is enabled, take a full database dump immediately after the resize operation completes. If multiple resizing operations are performed without taking transaction log dumps between each operation it may be necessary to manually sync the dSource with a new full database dump for Delphix to be able to continue ingesting source database dumps.

Target Host Requirements

- The operating system on the target environment must be the same as, or binary compatible with, the operating system on the source environment.
- As the Delphix Engine supports both NFSv3 and NFSv4 for mounting target host filesystems, the prerequisite packages that support NFSv3 or NFSv4 client communication are required for normal operation, and the required services to support NFS client communications (including file locking) must be running. This includes:
 - i. portmapper / rpcbind

- ii. status daemon (rpc.statd)
 - iii. NFS lock manager (rpc.lockd/lockmgr)
- The SAP ASE major version on the target environment must be the same as the version on the source environment. However, EBF/SP version on target environment can be different. If the target is used as a staging server, the ASE major version must be the same. The only caveat is for ASE 15.7 - where the source and staging patch levels must both be either any version below SP64 or both be any version above SP100.
- There must be an operating system user, such as delphix_os, that meets the following requirements:
 1. Set the PermitUserEnvironment configuration parameter to "yes" in the sshd_config file
 2. Add the variable to the user's .ssh/environment file
 3. Restart the SSH daemon
- The \$SYBASE environment variable is set for non-interactive shells (such as via the .bashrc configuration file). Set the variable as follows:
- To test this requirement:

```
ssh delphix_user@ase_hostname env | grep SYBASE
```

- Can login to the target host via Secure Shell (SSH)
- Can login to ASE instances using isql with LANG=C set
- Designating the Delphix operating system user's primary group to be the same as the ASE instance's means the file system permissions can be more restrictive and is a better security practice than granting world read access to the toolkit or the backup files. If the target host is used to host the staging databases, consider the following:
 - If you don't add the Delphix operating system user to the ASE instance owner's group, greater permissions will need to be **granted to the backup files to ensure read access to the dumps and/or transaction logs**. Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment).
 - If you sync Delphix with a dSource by asking ASE to create a new backup, the ASE instance owner will need **write permission to the toolkit** (or the mount point if you use the CLI to specify a directory other than the toolkit). Delphix will issue the "DUMP DATABASE" command to write to staging database's "temp" directory which is mounted on the staging host.
 - Has **write permission for the mount-point directory** (by default the toolkit directory but can be a separate mount point specified in the command line interface).
- The following permissions are usually granted via sudo authorization of the commands. See [Sudo Privilege Requirements for SAP ASE Environments](#) for further explanation of this requirement, and [Sudo File Configuration Examples for SAP ASE Environments](#) for example file configurations.
 1. Permission to run **mount** and **umount** as super-user.
 2. On Solaris, permission to run **pargs** on Solaris
 3. On AIX, permission to run the **nfso** command as super-user.
 4. (Optional) On AIX and Linux, permission to run **ps** as super-user.
 5. Disable **tty** for the delphix_os user for mount and umount.
- There must be a directory on the source host where you can install the Delphix platform toolkit, for example: /var/opt/delphix/Toolkit
 - The delphix_os user must own the directory
 - The directory must have permissions 0770, for example, -rwxrwx--. However, you can also use more permissive settings.
 - The directory should have 1GB of available storage
 - Avoid using the home directory of the delphix_os user
 - If you intend to use the LogSync feature, it is recommended to make the toolkit directory as short as possible to keep the full path to the transaction log file names under ASE's 127 character limit. For example, create the toolkit directory as /tk. Alternatively, link the dSource using the command line

interface and specify the "[mountBase](#)" parameter to mount the staging database's devices under a directory following your own naming convention.

Target Database Requirements

- There must be a database user, such as `delphix_db`, with the **sa_role** on each instance on the target environment
- The database user such as `delphix_db` for any staging instances must also have the **sybase_ts_role**
- If the target host will be used as a staging target environment, at least one of the following two options must be configured:
- You must use **sp_addserver** to add the staging ASE instance's Backup Server to **syservers** on the source ASE instance (so that remote database dump/load works)
OR
- Full and transaction dump files from the source database must be available locally to the staging database (over NFS, replication, scp, etc.)

Specific ASE Tuning Recommendations

- In ASE 15.7 SP60 and higher, there is a configuration parameter named "**enable large pool for load**". ASE automatically tunes the caches for recovery on reboot, but does not do this for load database and load transaction recovery, since it could impact other databases on a production server. Delphix recommends that the staging ASE instance be separate from the ASE production instance and ASE instances hosting VDBs so enabling this parameter should have **a beneficial impact on performance for the ASE staging instance**.
- Specify the "**relaxed**" strategy for replacing the cache in the default data cache (**sp_cacheconfig 'default data cache', relaxed**) in the ASE staging instance. Since the mount/unmount process invalidates the pages anyway, the page chain is really just unneeded overhead on the staging server.
- Staging and target ASE instances should have disk mirroring disabled.
 - sp_configure "disable disk mirroring"** – run value should be 1, which is the default. If it is 0, change it using
 - sp_configure "disable disk mirroring", 1** – this parameter is static so the ASE instance must be restarted for this change to take effect.
- Delphix will mirror the number of devices used on the source database for the staging database (dSource) and each VDB created from that source database. The number of devices parameter should be scaled appropriately based on the max number of virtual databases that will be provisioned to the ASE instance. This parameter can be changed using: **sp_configure "number of devices", <max devices>**.

i ASE 15.7.0 SP100 and later releases support the shrink command. In some cases, Delphix must increase the number of devices used for databases if this command is used. Delphix creates a minimum of the same number of devices as the source database for the staging database (dSource) and each VDB and will add more devices for every 4TB of fragment holes. See SAP ASE issue [CR#799273](#) for additional details.

To support multiple VDBs and the staging databases, you may need to increase the parameter number of **alarms**.

i Delphix uses ASE operations that use alarm structures such as **MOUNT** and **UNMOUNT**. The number of alarms limits the number of these operations which can be run concurrently. Various ASE instance failures can occur if the available alarm structures are exhausted. The amount of memory consumed by increasing

the number of alarm structures is small. Delphix recommends that the **number of alarms** value is increased to at least 4096.

Adding SAP ASE Source and Target Environments

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. From the **Actions (...)** menu select **Add Environment**.
5. In the **Host and Server** tab window, select **Unix/Linux**.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter a **Name** for the environment.
9. Enter the **Host IP** address.
10. For NFS Addresses (Optional):
Enter one or more comma-separated **IP Address/Hostname**

 If specified, Delphix Engine only allows NFS requests (mount, etc) originated from IP Addresses specified for the host.

11. Enter the **SSH** port.
The default value is **22**.
12. Select a **Login Type**.
 - Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For **Password Login**, click **Verify Credentials** to test the username and password.
14. Enter a **Toolkit Path**.
The toolkit directory stores scripts used for Delphix Engine operations. It must have a persistent working directory rather than a temporary one. The toolkit directory will have a separate subdirectory for each database instance. The toolkit path must have 0770 permissions.
15. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.

16. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
17. Click the **Discover SAP ASE** checkbox.
18. Click **Next**.
19. In the Summary, tab confirm your selections.
20. Click **Submit**.

ASE Manual Discovery

When an environment is added, Delphix discovers your ASE instances. Manual discovery allows users to add instances that were not automatically discovered. This feature is currently only supported via the CLI.

This topic describes how to use CLI commands to manually add ASE repositories to an SAP ASE environment. Discovery is the process by which the Delphix platform identifies data sources and data dependencies on a remote environment. ASE repository discovery is done automatically when an environment is added to the Delphix platform or when an already added environment is refreshed. In some cases, automatic discovery does not discover all of the repositories in an SAP ASE environment. These repositories may be added using manual discovery.

Unlike automatically discovered instances, manually discovered instances are not automatically deleted if the environment is refreshed when the instance isn't running. Manually discovered instances are not updated during an environment refresh either. So for example, if you upgrade ASE to a new version or change the listener port, you must manually update the repository.

To manually discover an ASE repository you will need to:

- Add the environment to Delphix
- Use CLI to manually discover a repository

Creating an ASE Environment

Please refer to [Adding an SAP ASE Environment](#) for detailed steps.

Manually discovering a repository

 In the following example, we are using **sc-dev3.dc2** as our example environment.

1. Log into CLI and cd to repository menu:

```
$ ssh delphix_admin@sc-dev3.dc2
Password:
sc-dev3.dc2> cd repository
sc-dev3.dc2 repository>
```

2. Add (manually discover) an ASE repository instance:

 The values used in the following code block are specific to the example instance we are adding.

```
sc-dev3.dc2 repository> create
sc-dev3.dc2 repository create *> ls
```

```

Properties
  type: ASEInstance
  credentials: (unset)
  dbUser: (unset)
  environment: (required)
  installationPath: (required)
  instanceName: (required)
  instanceOwner: (required)
  ports: (required)
  version: (unset)
sc-dev3.dc2 repository create *> set credentials.password=sybase
sc-dev3.dc2 repository create *> set dbUser=sa
sc-dev3.dc2 repository create *> set environment=sc-rhel64-sybase-ase-0
sc-dev3.dc2 repository create *> set installationPath=/opt/sybase/15-7
sc-dev3.dc2 repository create *> set instanceName=ASE1570_S1
sc-dev3.dc2 repository create *> set instanceOwner=sybase
sc-dev3.dc2 repository create *> set ports=5100
sc-dev3.dc2 repository create *> ls
Properties
  type: ASEInstance
  credentials:
    type: PasswordCredential (*)
    password: ***** (*)
  dbUser: sa (*)
  environment: sc-rhel64-sybase-ase-0 (*)
  installationPath: /opt/sybase/15-7 (*)
  instanceName: ASE1570_S1 (*)
  instanceOwner: sybase (*)
  ports: 5100 (*)
  version: (unset)
sc-dev3.dc2 repository create *> commit
`ASE_INSTANCE-22
sc-dev3.dc2 repository>

```

Updating a repository

Adding onto the above, the following example illustrates updating an ASE instance's version after upgrading ASE:

```

sc-dev3.dc2> repository
sc-dev3.dc2 repository> select ASE1570_S1
sc-dev3.dc2 repository 'ASE1570_S1'> update
sc-dev3.dc2 repository 'ASE1570_S1' update *> set version="15.7 SP138"
sc-dev3.dc2 repository 'ASE1570_S1' update *> ls
Properties
  type: ASEInstance
  credentials:
    type: PasswordCredential
    password: *****
  dbUser: sa
  installationPath: /opt/sybase/15-7
  instanceOwner: sybase
  linkingEnabled: true

```

```
ports: 5100
provisioningEnabled: true
servicePrincipalName: (unset)
staging: false
version: 15.7 SP138 (*)
sc-dev3.dc2 repository 'ASE1570_S1' update *> commit
```

 Be careful when setting the version string. Make sure that it matches output as displayed by the "select @@version" query all the way out to the patch level (PL). For example "15.7 SP138" or "16.0 SP02 PL01".

Enable Linking and Provisioning for SAP ASE Environments

This topic describes how to enable and disable provisioning and linking for SAP ASE databases.

Before a database can be used as a dSource, you must first make sure that you have enabled linking to its SAP ASE instance. Similarly, before you can provision a VDB to a target database, you must make sure that you have enabled provisioning to its SAP ASE instance.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **Databases** tab.
5. Click the **pencil** icon located next to the database **Installation Details**.
6. Select the **Allow Provisioning** checkbox to enable provisioning, deselect the checkbox to disable provisioning.
7. Click **show details** for the database.
8. Slide the button next to **Allow Linking** to **On** or **Off** to enable or disable linking.

Linking an SAP ASE Data Source

The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools. For an overview of all dSource related actions, please visit [Managing Data Sources](#).

Delphix Virtualization for SAP ASE databases leverages backup-based ingestion, which means that Delphix will look for, or sometimes initiate the creation of, a backup through your SAP ASE backup server. From there, the backup is restored on a staging server and the staging copy is then ingested into Delphix. See [Delphix Architecture with SAP ASE](#) for more information.

When linking a dSource from an SAP ASE source database, Delphix offers several different methods of capturing backup information:

- ASE Managed Backups, where the SAP ASE source database schedules and initiates backups. This method supports various backup types which include:
 - Full backups
 - Transaction log backups (with LogSync disabled)
 - Transaction log backups (with LogSync enabled)
- Delphix Managed Backups, where the Delphix Engine schedules and initiates the backups from the source database, and captures them.

ASE Managed Backups

Further contextual information on the various backup types (listed above):

- Full Backups - A snapshot will be created on the Delphix Timeflow for each Full backup.

- Transaction log backups (with LogSync disabled) - A snapshot will be created on the Delphix Timeflow for each transaction log backup.

Transaction Logs

Transaction logs are not collected if:

- a) there are gaps in the sequence of log backups (a break in the “log chain”).
 - b) the available log backups do not include any changes since the last successful Delphix snapshot.
- Transaction log backups (with LogSync enabled) - A snapshot will be created on the Delphix Timeflow for each transaction log backup. In addition, point-in-time provisioning will be an available option if you would like to provision from any point in between snapshots.

Log Files

Log files consume additional space on the Delphix Engine and are managed according to the defined retention policy for logs.

Delphix Managed Backups

When the checkbox for Delphix Managed Backups is selected, the Delphix Engine will initiate a full backup of the source database for the initial load of the dSource. Thereafter, the Delphix Engine will initiate full backups of the source database using the schedule specified by the selected SnapSync Policy. If you select the None policy, the Delphix Engine will not automatically initiate a full backup, but you can initiate them manually using the snapshot (camera) icon.

Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment)

Procedure

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.

 Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment).

4. In the **Add dSource** wizard, select the source environment with the correct environment-based user.
5. Enter your login credentials for the source database and click **Next**.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed, as described in [Data Management Settings for ASE Data Sources](#).
8. From the **Data Management** tab under the **Initial Load** option, select and enter any **additional settings** needed. There are three options for the initial load of the dSource:
 - a. If the source ASE instance resides on the same server as the staging ASE instance, the staging database's NFS mounted "temp" directory will be present for the source database to write to in response to the "DUMP DATABASE" command.
 - b. If the source and staging ASE instances are configured to allow remote access to the backup servers and the dSource is linked using the **Remote Server** option as described below. Delphix will then issue the "DUMP DATABASE" command and append the "AT <staging_backup_server_name>" clause so that the dump is written to the staging backup server.

1. (Recommended) **New Full Backup** - Lets Delphix create a new full backup file and load it. Note - that when Delphix creates the backup, it is moved to Delphix's NFS-mounted storage located on the stage host rather than. The backup will be located in the "temp" directory and will be deleted once the Delphix Engine has restored the backup and created a dSource from the restored staging database. This means that this option will work under two scenarios:
2. Most Recent **Existing Full Backup** – Find the most recent existing full backup file in the Backup Location and load it.

 If Dump History is not active on the Source Database: Choosing this option can delay completion of the dSource link as Delphix attempts to find and catalog every single backup listed in the source database's backup server log file.

3. Specific **Existing Full Backup** – Specify which backup files in the Backup Location you want to load. Choosing this option is much faster because Delphix will skip directly to loading the desired backup and only start to search for and catalog backups in the background after the linking of the source database has completed.

 When using a dump taken with the deprecated compression syntax, select the Specific Existing Full Backup option for Initial Load and, for each stripe, type compress::`<file name>` into the text box.

- Select the Staging environment and ASE instance that will be used to manage the staging database used for validated sync of the dSource.
- Select any policies for the new dSource.
- Click **Next**, then specify any pre-hook and post-hook scripts.
- Review the dSource Configuration and Data Management information, and click **Submit** to begin provisioning the VDB.

Provisioning an SAP ASE VDB

Procedure

1. In the Datasets panel on the left-hand side, click the group containing the dSource or VDB from which you want to provision.
2. Click the **Timeflow** tab.
3. Select a snapshot or open LogSync timeline to provision by a specific log or point in time.
 - Find more detail about initial provisioning options in the section 'Provisioning by Snapshot or LogSync' below.
 - Click to open the Provision VDB wizard, and select a compatible Target Environment for the new ASE VDB
 - Review the information presented for Target Configuration and edit as necessary.
 - Select a Snapshot Policy for the VDB.
 - (Optional) - Selective Data Distribution - After policies, there is a masking option.
 - Enter any operations that should be run in the Hooks page. These scripts can be managed after provision in the VDB's configuration page.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the Status tab, or by selecting Manage/Dashboards and viewing the Job History panel. Alternatively, you could see this in the Actions Sidebar. When provisioning is complete, the VDB will be included in the group you designated and listed in the Datasets panel. If you select the VDB in the Datasets panel and click the Configuration tab, you can view information about the database and its Data Management settings.

Provisioning by Snapshot or LogSync

When provisioning by Snapshot, you can provision to the start of any particular snapshot by time.

Provisioning By Snapshot/Time	Description
Provisioning By Snapshot	You can provision by using a Snapshot. In that case, a new VDB will be provisioned to the database state as of the Snapshot.
Provision by Time	If you have enabled Log Sync, you can provision a new database to a point in time. You can select a snapshot and then using time entry fields, specify a Point in Time. Delphix will use the selected snapshot to restore the VDB and use the log files to roll forward the VDB to the selected time.

Configuration Settings for ASE Virtual Databases

Each VDB has its own data management settings, found during the provisioning workflow as well as in the configuration page for that VDB. When you create a SAP ASE VDB, Delphix copies most configuration settings from the dSource and uses them to create the VDB. However, you can customize these with the following settings:

Setting	Explanation
Recovery Model	The current recovery model of the source database. This field will auto-populate with information from the dSource.
Auto VDB Restart	Enabling this option will automatically restart this VDB whenever its target host is rebooted.

Automatic VDB Restart on Target Server After Reboot

The Delphix platform now automatically detects whether a target server has been rebooted, and proactively restarts any VDB on that server that was previously up and running. This is independent of the data platform. It is done as if you realized a target server was restarted and issued a start command from the Delphix platform. This feature is compatible with Self-Service ordering dependencies and is limited to non-clustered VDBs.

To enable automatic restart, complete the following steps:

- When provisioning a new VDB in the VDB Provisioning wizard, check the **Auto VDB Restart** box.

Once the VDB has been provisioned, you will be able to turn **Automatic VDB Restart** on.

1. In the **Datasets** panel, select the VDB.
2. Select the **Configuration** tab.
3. Select **Source** sub-tab.
4. Select **Database edit**.

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.

2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision new VDBs from that masked VDB to quickly provide safe data to development and QA teams.

Related Topic

- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Provisioning and Managing VDBs from SAP ASE](#)

SAP ASE Support and Requirements

In order to begin using SAP ASE environments with Delphix, you will need to configure the source and target hosts with the requirements described in this section.

This section covers the following topics:

- [SAP ASE Support Matrix](#)
- [Requirements for SAP ASE Environments and Databases](#)
- [Network and Connectivity Requirements for SAP ASE Environments](#)
- [Sudo Privilege Requirements for SAP ASE Environments](#)
- [Sudo File Configuration Examples for SAP ASE Environments](#)

SAP ASE Support Matrix

Supported Operating Systems and Database Versions for SAP ASE

✔ Source and Target OS and DBMS Compatibility

The source and target must be running the same Operating System combination. The Operating System platform must be the same between the source and target, even when the operating system version may differ. For example, if the source is running Red Hat Enterprise Linux 6.2 x86_64 then the target could be running Red Hat Enterprise Linux 6.4 x86_64, but not Solaris 10 SPARC.

The source and target can be running the same DBMS version or the target ASE instance can have a higher DBMS version, (although users can run different patch/sp levels) in order to successfully provision a VDB to the target. For example, if the source is running SAP ASE 15.7 SP140, the target can be running ASE 15.7 SP138 or ASE 16.0 SP03.

If the target is used as a staging server, the source and target must be running the same DBMS combination, although users can run different patch/sp levels (prior to Delphix 5.2.5.0, the staging server had to match the source server's version down to the SP level). For example, if the source is running SAP ASE 16.0 SP03 PL03, the target can be running ASE 16.0 SP03 PL06. However, for ASE 15.7 the versions of the source and target should both be in the 15.7 SP100 to 15.7 SP140 or 15.7 ESD#1 to 15.7 SP64 ranges. There can be no mixing of versions between these two ranges.

Delphix 5.2.5.0 and higher relaxes Delphix's version checking of the staging instance. ASE itself may or may not allow the **MOUNT** command to work between versions. Be sure to check that the **UNMOUNT** and **MOUNT** commands work between your desired ASE versions.

So for example, if we upgrade the ASE staging instance to ASE 15.7 SP140 but we keep the ASE instance hosting the VDBs at ASE 15.7 SP135, we would want to make sure that we can **UNMOUNT** a database from ASE 15.7 SP140 and then **MOUNT** it on ASE 15.7 SP135:

```
isql -Usa -Psybase -SASE157SP140
1> UNMOUNT DATABASE testdb to "/tmp/manifest"
2> GO

isql -Usa -Psybase -S ASE157SP135
1> MOUNT DATABASE testdb from "/tmp/manifest"
2> go
Msg 14580, Level 16, State 1:
Server 'ASE157SP135', Line 1:
You cannot mount the database(s) because at least one database contains
functionality that is available only on the server on which it originated.
```

In the above example, ASE would not allow the database "testdb" to be mounted on the older release of ASE.



- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- 64-bit OS support only

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

 Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
RHEL 5.5				
RHEL 5.6				
RHEL 5.7				
RHEL 5.8				
RHEL 5.9				
RHEL 5.10				
RHEL 5.11				
RHEL 6.0				
RHEL 6.1				
RHEL 6.2				

RHEL 6.3	✓	✓	✓	✗
RHEL 6.4	✓	✓	✓	✗
RHEL 6.5	✓	✓	✓	✓
RHEL 6.6	✓	✓	✓	✓
RHEL 6.7	✓	✓	✓	✓
RHEL 6.8	✓	✓	✓	✓
RHEL 6.9	✓	✓	✓	✓
RHEL 6.10	✓	✓	✓	✓
RHEL 7.0	N/A	N/A	✓	✓
RHEL 7.1	N/A	N/A	✓	✓
RHEL 7.2	N/A	N/A	✓	✓
RHEL 7.3	N/A	N/A	✓	✓
RHEL 7.4	N/A	N/A	✓	✓
RHEL 7.5	N/A	N/A	✓	✓
RHEL 7.6	N/A	N/A	✓	✓
RHEL 7.7	N/A	N/A	✓	✓
RHEL 7.8	N/A	N/A	✓ 6.0.3+	✓ 6.0.3+

RHEL 7.9	N/A	N/A	✓ 6.0.4+	✓ 6.0.4+
RHEL 8.0	N/A	N/A	N/A	✓
RHEL 8.1	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.2	N/A	N/A	N/A	✓ 6.0.4+
RHEL 8.3	N/A	N/A	N/A	✓ 6.0.7+

SUSE Linux Enterprise Server (SLES)

 Support for SLES 15 requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Supported OS Version	Supported DBMS Version			
	15.0.3	15.5	15.7	16
SLES 11	✓	✗	✗	✗
SLES 11 SP1	✓	✗	✗	✗
SLES 11 SP2	✓	✗	✗	✗
SLES 11 SP3	✓	✗	✗	✗
SLES 11 SP4	✓	✗	✗	✗
SLES 12	✗	✗	✗	✗
SLES 12 SP1	✗	✗	✗	✗
SLES 12 SP2	✗	✗	✗	✗

SLES 12 SP3	✗	✗	✗	✗
SLES 12 SP4	✗	✗	✗	✓ 6.0.4+

Solaris SPARC

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
Solaris 10 U9	✓	✓	✓	✗
Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	✓	✓	✓	✗
Solaris 11	✗	✗	✗	✗
Solaris 11 U1	✗	✗	✗	✗
Solaris 11 U2	✗	✗	✗	✗
Solaris 11 U3	✗	✗	✓	✗
Solaris 11 U4	✗	✗	✓ 6.0.1.0+	✓ 6.0.1.0+

Solaris x86

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
Solaris 10 U9	✓	✓	✓	✗
Solaris 10 U10	✓	✓	✓	✗
Solaris 10 U11	N/A	N/A	✓	✗

Solaris 11	N/A	N/A	✓	✓
Solaris 11 U1	N/A	N/A	✓	✓
Solaris 11 U2	N/A	N/A	✓	✓
Solaris 11 U3	N/A	N/A	✓	✓
Solaris 11 U4	N/A	N/A	✗	✗

Hewlett Packard Unix (HP-UX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
HP-UX 11.31	✗	✗	✗	✗

Advanced Interactive eXecutive (AIX)

	Supported DBMS Version			
Supported OS Version	15.0.3	15.5	15.7	16
AIX 6.1	✓	✗	✓	✗
AIX 7.1	✗	✗	✓	✓ 6.0.1.0+
AIX 7.2	✗	✗	✓	✓

Related Topics

- [SAP ASE Support and Requirements](#)

Requirements for SAP ASE Environments and Databases

In order to begin using SAP ASE environments with Delphix, you will need to configure the source and target hosts with the requirements described on this page.

SAP ASE Host Requirements

These requirements are applicable to source, staging as well as target environments. However, target environments have additional requirements which are detailed in the **Target Host Requirements** section below.

On each host with SAP ASE, there must be an operating system user (e.g. delphix_os) configured to the required specifications for Delphix.

Host Requirement	Explanation
Profile and privileges should be the same as the ASE user (e.g. Sybase) on the host.	For example, delphix_os should have the same environment variable settings (e.g. \$PATH), umask, and ulimit settings, as the user sybase.
Delphix operating system user's primary group should be the same as SAP ASE instance's group.	To perform Sync/Validated Sync operations on the source database, Delphix reads a number of files like SAP ASE Backup server log file, ASE Dump History file, and backup files. If Delphix OS user does not have the same group as of the SAP ASE instance, these files will need "everyone" read permissions, which is less restrictive.
Delphix Engine Toolkit Directory on Source Host (e.g. /var/opt/delphix/toolkit) <ul style="list-style-type: none"> The delphix_os user must own the directory The directory must have permissions 0770, for example, -rwxrwx---. However, you can also use more permissive settings. The directory should have 256MB of available storage 	<ul style="list-style-type: none"> If you sync Delphix with a dSource by asking SAP ASE to create a new backup, the SAP ASE instance owner will need write permissions to the toolkit (or the mount point if you use the CLI to specify a directory other than the toolkit). Delphix will issue the "DUMP DATABASE" command to write to the staging database's "temp" directory which is mounted on the staging host. Has write permission for the mount-point directory (by default the toolkit directory but can be a separate mount point specified in the command line interface)
Backup Server Log File Permissions	<ul style="list-style-type: none"> In the case that the Delphix OS user is not in the same group as the SAP ASE instance user, you need to change permissions to SAP ASE Backup server log file from "rw-r---" to "rw-r-r--". The "rw-r---" to "rw-r-r--" permissions change is also required if Dump History files are in use.
The Delphix Engine must be able to make an SSH connection to the source host (typically port 22).	

SAP ASE Database Requirements

Delphix for SAP ASE requires a specific login to run discovery and linking activities for instances and databases. Delphix requires specific user permissions for database discovery and linking (detailed below).

 Discovery and Linking can be separated into separate users if desired.

 If you are using an ASE backup server instance shared with multiple ASE data server instances, the use of the dump history file feature is mandatory. This should be enabled on all instances and the dump history file in a location accessible by the Delphix OS user account from the Staging host.

SAP ASE Source Database Requirements

The discovery database user (delphix_disc for example) must have SELECT privileges on the following tables for each SAP ASE instance on the source host:

- Sysdatabases
- Sysservers
- Syslisteners
- Sysconfigures
- Syscurconfigs

Database Linking

Permissions for linking environments can be included under the “discovery” user or can be a separate, dedicated role - e.g. delphix_link. If separate, the “linking” user must be specified when linking each dSource (delphix_link for example) that has SELECT privileges on the above tables.

If you will select “No new Full Backup” when linking, this user must also have privileges to take a new full database dump of the source database. For more information about linking options, see [Linking an SAP ASE Data Source](#).

 The link database user can be different for each instance and database on the source host.

Sample Script to create delphix_link on Linux

```
note: run as sa

sp_addlogin delphix_link, "StrongPassword"

go

sp_adduser delphix_link

go

grant select on sysdatabases to delphix_link

go
```

```
grant select on sysservers to delphix_link
go
grant select on syslisteners to delphix_link
go
```

Requirements when resizing Source Databases

For SAP ASE on Delphix, transaction logs are helpful for maintaining a Timeflow via logsync. However, not all customers will want to maintain these logs for bandwidth purposes. The following section will help guide users on requirements to maintain timeflows when it comes to resizing databases with and without (` trunc log on chkpt`).

Active SAP ASE Feature on Source Host	Requirement
(` trunc log on chkpt`) is disabled	take a transaction log dump immediately after the resize operation completes.
(` trunc log on chkpt`) is enabled	take a full database dump immediately after the resize operation completes.

 If multiple resizing operations are performed without taking transaction log dumps between each operation it may be necessary to manually sync the source database dumps.

Additional Target Host Requirements

This section describes the user privileges, and environment discovery requirements, that are required for SAP ASE target hosts and databases, collectively referred to as target environments.

Target Host Requirement	Explanation
The operating system on the target environment must be the same as, or binary compatible with, the operating system on the source environment	
As the Delphix Engine uses NFSv3 for mounting target host filesystems, the prerequisite packages to support NFSv3 client communication are required for normal operation, and the required services to support NFS client communications (including file locking) must be running This includes: <ul style="list-style-type: none"> • portmapper / rpcbind • status daemon (rpc.statd) • NFS lock manager (rpc.lockd/lockmgr) 	The Delphix Engine uses NFSv3 as a defacto mounting standard for target host file systems.

SAP ASE Source/Target Environment Version Compatibility	<ul style="list-style-type: none"> The major/minor/macro version of the target ASE instance should be the same or higher than the version of the source ASE instance. Delphix supports the ASE databases to be restored from lower to higher version or within different patch/sp levels as long as SAP ASE supports it. If the target environment is used as staging then the SAP ASE version on the target environment must be the same as the version on the source environment. However, EBF/SP version on the target environment can be different.
Create OS User (e.g. delphix_os)	There must be an operating system user, such as delphix_os, that meets the following requirements:
OS User must have - The \$SYBASE environment variable set for non-interactive shells (such as via the .bashrc configuration file)	<p>Additional Settings for the \$SYBASE environment variable:</p> <ul style="list-style-type: none"> Set the PermitUserEnvironment configuration parameter to "yes" in the sshd_config file Add the variable to the user's .ssh/environment file Restart the SSH daemon
OS User must have appropriate access to Target Hosts/Instances	<ul style="list-style-type: none"> Can login to the target host via Secure Shell (SSH) Can login to ASE instances using isql with LANG=C set
OS User's Primary Group must be the same as ASE Instance's Primary Group	<ul style="list-style-type: none"> If not, the file system permissions can be more restrictive. This also is a better security practice than granting world read access to the toolkit or the backup files.
<p>Permissions granted via sudo authorization of the commands:</p> <ul style="list-style-type: none"> Disable tty for the delphix_os user for mount and umount Permission to run mount and umount as super-user. On Solaris, permission to run pargs on Solaris On AIX, permission to run the nfsd command as super-user. (Optional) On AIX and Linux, permission to run ps as super-user. 	See Sudo Privilege Requirements for SAP ASE Environments for further explanation of this requirement, and Sudo File Configuration Examples for SAP ASE Environments for example file configurations.

<p>There must be a database user, such as delphix_db, with the sa_role on each instance on the target environment</p>	
<p>The database user such as delphix_db for any staging instances must also have the sybase_ts_role</p>	
<p>There must be a directory on the target environment where you can install the Delphix Engine toolkit, for example, /var/opt/delphix/Toolkit.</p>	<ul style="list-style-type: none"> • The delphix_os user must own the directory • The directory must have permissions 0770, for example, -rwxrwx--. However, you can also use more permissive settings. • The directory should have 1GB of available storage • Avoid using the home directory of the delphix_os user • If you intend to use the LogSync feature, it is recommended to make the toolkit directory as short as possible to keep the full path to the transaction log file names under ASE's 127 character limit. For example, create the toolkit directory as /tk. Alternatively, link the dSource using the command line interface and specify the "mountBase" parameter to mount the staging database's devices under a directory following your own naming convention.
<p>(Optional - If the target host will be used as a staging target environment)</p> <p>Must have at least one of the following two options configured:</p> <ul style="list-style-type: none"> • use sp_addserver to add the staging ASE instance's Backup Server to syssservers on the source ASE instance (so that remote database dump/load works) <p>OR</p> <ul style="list-style-type: none"> • Full and transaction dump files from the source database must be available locally to the staging database (over NFS, replication, scp, etc.) 	<p>(see Managing SAP ASE Environments)</p>

Special Considerations - OS User Settings (If Target Host is hosting Staging DBs)

Target Host Requirement	Explanation
-------------------------	-------------

If you don't add the Delphix operating system user to the SAP ASE instance owner's group, greater permissions will need to be granted to the backup files to ensure read access to the dumps and/or transaction logs.	Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment).
ASE Instance Owner Write Permissions to the toolkit (or mount point if you use the CLI to specify a directory other than the toolkit)	If you sync Delphix with a dSource by asking ASE to create a new backup, the SAP ASE instance owner will need write permission to the toolkit (or the mount point if you use the CLI to specify a directory other than the toolkit). Delphix will issue the "DUMP DATABASE" command to write to the staging database's "temp" directory which is mounted on the staging host.
Has write permission for the mount-point directory	The mount-point directory defaults to the toolkit directory, which has write permissions. However, the mount-point directory can be customized, it only needs to have write permissions.

Additional ASE Tuning Recommendations

- In SAP ASE 15.7 SP60 and higher, there is a configuration parameter named "enable large pool for load". SAP ASE automatically tunes the caches for recovery on reboot but does not do this for load database and load transaction recovery, since it could impact other databases on a production server. Delphix recommends that the staging SAP ASE instance be separate from the SAP ASE production instance and SAP ASE instances hosting VDBs so enabling this parameter should have *a beneficial impact on performance for the SAP ASE staging instance*.
- Specify the "relaxed" strategy for replacing the cache in the default data cache (sp_cacheconfig 'default data cache', relaxed) in the SAP ASE staging instance. Since the mount/unmount process invalidates the pages anyway, the page chain is really just unneeded overhead on the staging server.
- Staging and target SAP ASE instances should have disk mirroring disabled.
sp_configure "disable disk mirroring" – run value should be 1, which is the default. If it is 0, change it using sp_configure "disable disk mirroring", 1 – this parameter is static so the ASE instance must be restarted for this change to take effect.
- Delphix will mirror the number of devices used on the source database for the staging database (dSource) and each VDB created from that source database. The number of devices parameter should be scaled appropriately based on the max number of virtual databases that will be provisioned to the SAP ASE instance. This parameter can be changed using: sp_configure "number of devices", <max devices>. SAP ASE 15.7.0 SP100 and later releases support the shrink command. In some cases, Delphix must increase the number of devices used for databases if this command is used. Delphix creates a minimum of the same number of devices as the source database for the staging database (dSource) and each VDB and will add more devices for every 4TB of fragment holes. See SAP ASE issue [CR#799273](#) for additional details.

To support multiple VDBs and the staging databases, you may need to increase the parameter number of alarms.

i Delphix uses SAP ASE operations that use alarm structures such as MOUNT and UNMOUNT. The number of alarms limits the number of these operations which can be run concurrently. Various SAP ASE instance failures can occur if the available alarm structures are exhausted. The amount of memory consumed by

increasing the number of alarm structures is small. Delphix recommends that the number of alarms value is increased to at least 4096.

ASE Manual Discovery

When an SAP ASE environment is added Delphix automatically discovers your SAP ASE instances. Manual discovery allows users to add instances that were not automatically discovered. This feature is currently only supported via the CLI. For more information please refer to [Configuring ASE Manual Discovery](#).

Related Topics

- [SAP ASE Support and Requirements](#)

Network and Connectivity Requirements for SAP ASE Environments

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.

Protocol	Port Number	Use
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.


Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

Connection Requirements for SAP ASE Environments

- The Delphix Engine uses an **SSH** connection to each source environment and **SAP ASE client** connections to the SAP ASE instances on the source environment.
- The Delphix Engine uses an **SSH** connection to each target environment, **NFS** connections from each target environment to the Delphix Engine, and **SAP ASE JDBC** connections to the virtual databases on the target environment.

Port Allocation for SAP ASE Environments

The following diagram describes the port allocations for SAP ASE environments. It illustrates the ports that we recommend to be open from Delphix to remote services, to the Delphix Engine, and to the Target Environments.



Refer to [Managing SAP ASE Environments](#) for information on SAP ASE environments. The Delphix Engine makes use of the following network ports for SAP ASE dSources and VDBs:

Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	Configuration dependent	JDBC Connections to the SAP ASE instances on the source environments

Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
UDP	33434-33464	Traceroute from source and target database servers to the Delphix Engine (optional)

Protocol	Port Number	Use
TCP/UDP	111	Remote Procedure Call (RPC) port mapper used for NFSv3 mounts
TCP	2049	NFS client from target hosts to the Delphix Engine (NFSv3 and NFSv4)
TCP	1110	Network Status Monitor (NSM) client from target hosts to Delphix Engine
TCP	54043	Client mount daemon (NFSv3 only)
TCP	54044	Lock state notification service (NFSv3 only)
TCP	54045	Network Lock Manager (NLM) client from target hosts to Delphix Engine (NFSv3 only)

Port Allocation Between Source and Staging Target Environments

Protocol	Port Numbers	Use
TCP	Configuration dependent	SAP ASE Remote Backup Server protocol. Applies if linking using the New Full Backup option, or if linking with the Remote Backup Server option.

Port Allocation Between Staging Target Environments and Shared Backup Fileserver

Protocol	Port Numbers	Use
TCP/UDP	NFS and related port numbers: <ul style="list-style-type: none"> • Portmap (111) • NFS (2049) • Network Lock Manager (NLM) • Network Status Monitor (NSM) 	NFS mount point exported by an NFS shared backup fileserver. Applies if linking using the Local Backup Server option.

AppData Port Requirements

The use of AppData requires the following ports/protocols.

Two important notes about these specifications:

1. The next release of the Delphix Engine will significantly augment the port/protocol utilization of AppData. The upcoming-only requirements have been marked with a *.
2. AppData V2P uses RSYNC to export to the target. RSYNC between the target and Delphix Engine is not required for general virtualization usage. The V2P-only requirements have been marked with a ^.

From Source to Delphix Engine	From Delphix Engine to Source	From Target to Delphix Engine	From Delphix Engine to Target
RSYNC (TCP Port 873)	RSYNC (TCP Port 873)	DSP (Default TCP Port 8415)	DSP (Default TCP Port 8415)
DSP (Default TCP Port 8415)	SSH (TCP Port 22)	NFS	SSH (TCP Port 22)
*NFS	DSP (Default TCP Port 8415)	^RSYNC (TCP Port 873)	^RSYNC (TCP Port 873)

Related Topics

- [SAP ASE Support and Requirements](#)
- [Managing SAP ASE Environments and Hosts](#)
- [Configuring Replication](#)

Sudo Privilege Requirements for SAP ASE Environments

This topic describes the rationale behind specific `sudo` privilege requirements for virtualizing SAP ASE Databases.

Privilege	Sources	Targets	Rationale
<code>pargs</code>	Required on Solaris	Required on Solaris	Delphix attempts to call <code>pargs</code> to discover the arguments of the ASE processes. It needs the name of each running dataserver or backupserver process so that it can try to connect to the instances to gather further information during the discovery process.
<code>ps</code>	Optional on Linux, AIX	Optional on Linux, AIX	<p>Delphix attempts to call <code>ps</code> to discover the arguments of the ASE processes. It needs the name of each running dataserver or backupserver process so that it can try to connect to the instances to gather further information during the discovery process.</p> <p>Unlike Solaris, Delphix can usually determine the arguments without <code>sudo</code> privileges on Linux/AIX. But Delphix will attempt "<code>sudo ps</code>" before attempting a regular <code>ps</code> command, and this could cause locking of the delphix_os account. To avoid locking issues, you can grant <code>sudo ps</code> to delphix_os.</p>
<code>mount/umount</code>	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for superuser.

Privilege	Sources	Targets	Rationale
<code>nfso</code>	Not Required	Required on AIX	Delphix monitors NFS read and write sizes on an AIX target host. It uses the <code>nfso</code> command to query the sizes in order to optimize NFS performance for VDBs running on the target host. Only a superuser can issue the <code>nfso</code> command.

Default Mount Directory

By default, Delphix mounts the NFS directories for VDBs and staging databases under the toolkit directory. Sudo permissions should be granted to allow the mount/umount commands to execute under these directories unless the dSource is linked using the command-line interface (CLI) and a different NFS mount base is specified. Please refer to the Reference manual for more information on linking the dSource using the CLI and specifying the "`mountBase`" parameter.

Specify the NOPASSWD qualifier

It is required to specify the NOPASSWD qualifier within the "sudo" configuration file, as shown here: [Sudo File Configuration Examples for SAP ASE Environments](#). This ensures that the "sudo" command does not demand the entry of a password, even for the "display permissions" (i.e. "sudo -l") command. Delphix issues "sudo -l" in some scripts to detect if the operating system user has the correct sudo privileges. If it is unable to execute this command, some actions may fail and Delphix will raise an alert suggesting it does not have the correct sudo permissions. Restricting the execution of "sudo -l" by setting "listpw=always" in the "/etc/sudoers" file **when the Delphix operating system user is configured to use public key authentication** will cause the Delphix operating system user to be prompted for a password which will fail certain Delphix actions. Use a less restrictive setting for listpw than "always" when the Delphix operating system user is using public-key authentication.

SAP ASE and AppData Mount Options

AIX	<code>-o cio,rw,fg,hard,rsize=\$nfs_rsize,wsiz=\$nfs_wsize,nointr,timeo=600,proto=tcp,noacl</code>
HPUX	<code>-o rw,hard,rsize=1048576,wsiz=1048576,nointr,timeo=600,proto=tcp,suid</code>
Solaris	<code>-F nfs -o rw,fg,hard,rsize=1048576,wsiz=1048576,nointr,timeo=600,proto=tcp,suid,sec=sys</code>

For these platforms, depending on the NFS version used, additional options `vers=3` or `vers=4.x` is added (x varies depending on what that platform supports. e.g. `vers=4` or `vers=4.1`)

Linux (NFSv3)	<code>-t nfs -o rw,fg,hard,rsize=1048576,wsize=1048576,nointr,ti meo=600,tcp,noacl,vers=3</code>
Linux (NFSv4)	<code>-t nfs4 -o rw,fg,hard,rsize=1048576,wsize=1048576,nointr,ti meo=600,sec=sys,tcp,noacl</code>
(For some flavors of Linux and NFSv4.1, additional optional 'v4.1' is added)	

-  1. AppData plugins and toolkits have some additional mount options depending on the type of toolkit/plugin.
2. "port=2049" is added for all the platforms.

SAP ASE and AppData unmount options

"-f" in most cases. Certain cases, SAP ASE uses "-lf". (Lazy unmount option)

Mount and Unmount Options Subject to Change

Please note that the mount and unmount options listed above are subject to change. For example, if Delphix finds that a certain option improves performance, Delphix may add, remove or change options at anytime. Therefore, it is highly recommended to create the sudo profiles using wildcards that allow any number of options.

Related Topics

- [SAP ASE Support and Requirements](#)
- [Sudo File Configuration Examples for SAP ASE Environments](#)

Sudo File Configuration Examples for SAP ASE Environments

This topic provides sample `sudo` file privilege configurations for using the Delphix Engine with various operating systems and SAP ASE.

Configuring `sudo` Access on Solaris for SAP ASE Source and Target Environments

Sudo access to `pargs` on the Solaris operating system is required to discover the arguments of the ASE processes both source and target environments.

Example: Solaris `/etc/sudoers` entries for a Delphix Source for SAP ASE

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/usr/bin/pargs
```

On a Solaris target, `sudo` access to `mount` and `umount` is also required.

Example: Solaris `/etc/sudoers` entries for a Delphix Target for SAP ASE

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

User_Alias   DELPHIX_USER=delphix_os

Cmd_Alias   DELPHIX_CMDS= \
/usr/sbin/mount, \
/usr/sbin/umount, \
/usr/bin/pargs

DELPHIX_USER ALL=(ALL) NOPASSWD: DELPHIX_CMDS
```

Configuring `sudo` Access on Linux for SAP ASE Source and Target Environments

On a Linux target, sudo access to `mount` and `umount` is required.

Example: Linux `/etc/sudoers` file for a Delphix Target for SAP ASE

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

Defaults:delphix_os !requiretty
```

```
delphix_os ALL=NOPASSWD: \
/bin/mount, /bin/umount
```

Configuring `sudo` Access on AIX for SAP ASE Source and Target Environments

In addition to `sudo` access to the `mount` and `umount` commands on AIX target hosts, Delphix also requires `sudo` access to `nfso`. This is required on target hosts for the Delphix Engine to monitor the NFS read write sizes configured on the AIX system. Super-user access level is needed to run the `nfso` command.

Example: AIX `/etc/sudoers` File for a Delphix Target

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
/bin/mount, \
/bin/umount, \
/usr/sbin/nfso
```

Examples of Limiting `sudo` Access for the Delphix OS User

In situations where security requirements prohibit giving the Delphix user root privileges to `mount`, `umount`, `mkdir`, and `rmdir` on the global level, it is possible to configure the `sudoers` file to provide these privileges only on specific mount points or from specific Delphix Engines, as shown in these two examples.

i The Delphix Engine tests its ability to run the `mount` command using `sudo` on the target environment by issuing the `sudo mount` command with no arguments. Many of the examples shown in this topic do not allow that. This causes a warning during environment discovery and monitoring but otherwise does not cause a problem. If your VDB operations succeed, it is safe to ignore this warning. Similarly, the `ps` or `pargs` the command is used for target environment operations such as initial discovery and refresh. Some organizations configure the security on the target environments to monitor `sudo` failures and lockout the offending account after some threshold. In those situations, the failure of the `sudo` commands might cause the **delphix_os** account to become locked. One workaround for this situation is to increase the threshold for locking out the user account. Another option is to modify `/etc/sudoers` to permit the **delphix_os** user to run `ps` (`pargs`), `mkdir`, `rmdir`, and `mount` command without parameters.

⚠ Note that the following examples are for illustrative purposes and the `sudo` file configuration options are subject to change.

Example 1

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/sybase`.

Note that wildcards are allowed for the options on `mount` and `umount` because those commands expect a fixed number of arguments after the options. The option wildcard on the `mount` command also makes it possible to specify the file-system being mounted from the Delphix Engine.

However, wildcards are not acceptable on `mkdir` and `rmdir` because they can have any number of arguments after the options. For those commands, you must specify the exact options (`-p`, `-p -m 755`) used by the Delphix Engine.

Example /etc/sudoers File Configuration on the Target Environment for sudo Privileges on the VDB Mount Directory Only (Linux OS)

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount *          /sybase/*, \
/bin/mount "", \
/bin/umount *        /sybase/*, \
/bin/umount          /sybase/*, \
/bin/mkdir -p        /sybase/*, \
/bin/mkdir -p -m 755 /sybase/*, \
/bin/mkdir           /sybase/*, \
/bin/rmdir           /sybase/*, \
/bin/ps
```

Example /etc/sudoers File Configuration on the Source Environment to grant Super-User privileges when running PS

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: /bin/ps
```

Example 2

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/sybase`, restricts the mount commands to a specific Delphix Engine hostname and IP, and restricts user-specified options for the `umount` command.

This configuration is more secure, but there is a tradeoff with deployment simplicity. This approach would require a different sudo configuration for targets configured for different Delphix Engine.

Configuring the /etc/sudoers File on the Target Environment for Privileges on the VDB Mount Directory Only, and Allows Mounting only from a Single Server (Linux OS)

```
# Delphix issues sudo -l so we need to allow it via listpw. Never set it to always
when using public key authentication
Defaults      listpw=all

Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount      <delphix-server-name>* /sybase/*, \
/bin/mount *    <delphix-server-name>* /sybase/*, \
/bin/mount      <delphix-server-ip>* /sybase/*, \
/bin/mount *    <delphix-server-ip>* /sybase/*, \
/bin/mount "", \
/bin/umount     /sybase/*, \
/bin/umount *   /sybase/*, \
/bin/mkdir [*]  /sybase/*, \
/bin/mkdir      /sybase/*, \
/bin/mkdir -p   /sybase/*, \
/bin/mkdir -p -m 755 /sybase/*, \
/bin/rmdir      /sybase/*, \
/bin/ps
```

Related Topics

- [Sudo Privilege Requirements for SAP ASE Environments](#)
- [SAP ASE Support and Requirements](#)

Managing SAP ASE Environments and Hosts


This section describes the attributes of SAP ASE environments such as information for the Delphix Connector and covers the following topics:

- [Managing SAP ASE Environments Overview](#)
- [Adding an SAP ASE Environment](#)
- [Environment Attributes for Hosts with SAP ASE](#)
- [How to Discover SAP ASE Instances Which Use Multiple Network Handlers](#)
- [Changing the Host Name or IP Address of an SAP ASE Environment](#)
- [Using HostChecker to Validate SAP ASE Source and Target Environments](#)
- [Enabling Linking and Provisioning for SAP ASE Environments](#)

Managing SAP ASE Environments Overview

This topic describes the high-level process for adding SAP ASE environments, linking SAP ASE databases to the Delphix Engine, and provisioning virtual databases.

Block Diagram of Linking Architecture between SAP ASE Environments and the Delphix Engine



Environment Setup

SAP ASE dSources are backed by a staging database that runs on a target host, as shown in the diagram. There is no requirement for additional local storage on this host, as the storage is mounted over NFS from the Delphix Engine. At Delphix, we refer to the creation and maintenance of this staging database on the staging host as "validated sync," because it prepares the dSource data on the Delphix Engine for provisioning VDBs later on. After the Delphix Engine creates the staging database, it continuously monitors the source database for new transaction log dumps. When it detects a new transaction log dump, it loads that dump to the staging database. The result is a TimeFlow with consistent points from which you can provision a virtual database (VDB), and a faster provisioning process because there is no need for any database recovery during provisioning.

When you later provision a VDB, you can specify any environment as a target, including the environment that contains the staging database. However, for best performance, Delphix recommends that you choose a different

target environment. The target must have an operating system that is compatible with the one running on the validated host, as described in [Requirements for SAP ASE Environments and Databases](#).

Workflow and Tasks for SAP ASE Environments

1. Add the desired source environments as described in [Managing SAP ASE Environments](#).
2. Add the desired target environments as described in [Managing SAP ASE Environments](#).
3. Link the source database as described in [Linking Data Sources with SAP ASE](#).
4. Provision VDBs as described in [Provisioning an SAP ASE VDB](#).

Best Practice

If possible, it is recommended to use separate SAP ASE instances for staging databases (dSources) and VDBs on target environments. This allows better control on target environments based on the resources required for SAP ASE full/transaction log dump ingestion vs VDB performance.

Related Topics

- [SAP ASE Support and Requirements](#)
- [Managing SAP ASE Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Provisioning and Managing VDBs from SAP ASE](#)

Adding an SAP ASE Environment

Prerequisites

See [Requirements for SAP ASE Environments and Databases](#).

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. From the **Actions (...)** menu select **Add Environment**.
5. In the **Host and Server** tab window, select **Unix/Linux**.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter a **Name** for the environment.
9. Enter the **Host IP** address.
10. For NFS Addresses (Optional):
Enter one or more comma-separated **IP Address/Hostname**

 If specified, Delphix Engine only allows NFS requests (mount, etc) originated from IP Addresses specified for the host.

11. Enter the **SSH** port.
The default value is **22**.
12. Select a **Login Type**.
 - Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For **Password Login**, click **Verify Credentials** to test the username and password.
14. Enter a **Toolkit Path**.
The toolkit directory stores scripts used for Delphix Engine operations. It must have a persistent working

directory rather than a temporary one. The toolkit directory will have a separate subdirectory for each database instance. The toolkit path must have 0770 permissions.

15. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
16. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
17. Click the **Discover SAP ASE** checkbox.
18. Click **Next**.
19. In the Summary, tab confirm your selections.
20. Click **Submit**.

Post-Requisites

After you create the environment, you can view information about it by selecting **Manage > Environments** and then selecting the **environment name**.

Configuring a non-default Dump History File

Delphix, by default, uses the Dump History file that is configured for the source ASE instance. Delphix will query the source ASE instance to find this file name and use this file to find dump history. However, users can specify any other file on the source host to be used for querying the backup history. In that case, Delphix will not use the file configured for the source ASE instance but will use the file specified to get backup information. To specify a custom Dump History file, go to the **source environment > Databases**. Click the pencil button in front of **Installation Details** for the desired instance. Enter the fully qualified name of the Dump History file to be used in the text box next to **Dump History File**.

The screenshot displays the Delphix Management interface. On the left, the 'Environments' sidebar shows 'ASE_RH74SRC' (Unix Host - Enabled). The main area is titled 'Installations' and shows a dropdown for 'ASE157_SRC'. A modal window is open for 'Installation Details' for this instance. In this modal, the 'Dump History File' field is highlighted with a blue border and contains the path '/opt/sybase/15-7/sp139/instal'. Other fields in the modal include 'Version' (15.7 SP139), 'Ports' (5400), 'Page Size' (4,096), and 'Owner' (sybase). To the right of the modal, the 'Databases' section lists 'cozumel' and 'ibiza', both with 'Allow Linking Enabled' and 'Add dSource' buttons. Below these are two database entries: 'dx7s633tnzYtnB5lr6Mzx_cayman (dSource: cayman)' and 'cayman (dSource: cayman)'. The modal has 'X' and checkmark buttons at the bottom.

Delphix does not support rolled over Dump History Files.

Related Topics

- [Requirements for SAP ASE Environments and Databases](#)
- [Managing SAP ASE Environments and Hosts](#)

Environment Attributes for Hosts with SAP ASE

This topic describes the attributes of SAP ASE environments. Below you will see a section for common environment attributes shared by all types of environments as well as SAP ASE specific ones.

Procedure

1. From your Delphix Engine, click **Manage**.
2. Select **Environments**.
3. In the **Environments** panel, click the name of an environment to view its attributes.
4. Next to **Attributes**, click the Pencil icon to edit an attribute.

Common Environment Attributes

Attribute	Description
Environment Users	The users for that environment. These are the users who have permission to ssh into an environment or access the environment through the Delphix Connector. For more information on the environment user requirements, see the Requirements topics for specific data platforms.
Host Address	The IP address of the environment host.
DSP KeyStore Path	The path to the user-managed DSP Keystore.
DSP KeyStore Alias	The lowercase alias to use inside the user-managed DSP Keystore.
DSP KeyStore Password	The password for the user-managed DSP Keystore.
DSP TrustStore Path	The path to the user managed DSP truststore.
DSP TrustStore Password	The password for the user managed DSP truststore.
OS	The name of the host operating system.
Version	The version of the host operating system.

Attribute	Description
Release	The release of the host operating system.
Time Zone	The timezone of the host operating system.
Total RAM	The amount of RAM on the host machine.
Processor Type	The processor type of the host machine.
Traceroute	Traceroute info from target host to Delphix Engine.
Notes	Any other information you want to add about the environment.
Java Development Kit	The currently selected JDK kit will be shown.
Java Development Kit (JDK) Path	Location of the Java Development Kit (JDK) used for the host. Only specified if the feature to provide your own JDK is enabled, otherwise, the defaults are used per our Java Support Matrix .

SAP ASE Environment Attributes

Attribute	Description
Database User	User for Delphix to use for SAP ASE database operations.
DB Password	Credentials to use for the database user in this environment.

Related Topics

- [Managing SAP ASE Environments and Hosts](#)

How to Discover SAP ASE Instances Which Use Multiple Network Handlers

When adding a SAP ASE or Sybase ASE environment, the Delphix Engine may not be able to discover ASE instances that are not listening on the same network address that the engine used for discovery. In these cases, you may receive the following warning:

```
WARNING: Error during discovery for instance "INSTANCENAME" : "Failed to connect to instance "INSTANCENAME over JDBC.". Skipping discovery for instance "INSTANCENAME".
```

Troubleshooting

An ASE instance may have been configured to listen on other network interfaces for one of the following reasons:

- To support the use of clustering/failover technologies such as Veritas Cluster Server
- To support the concept of "virtual hostnames" or "virtual IPs" (VIPs) to provide abstraction for client applications
- To facilitate a previous database migration to new infrastructure
- Database administrators configured it manually for other reasons

For guidance and best practices for configuring the Delphix Engine to work with clustering or failover technologies, contact your Professional Services or Customer Success representative.

To verify whether an ASE instance is using multiple network interfaces, use the `sp_listener` stored procedure when connected to the instance using `isql` :

```
1> sp_listener status
2> go
proto      host                                     port      status
-----
tcp        secondary-hostname-or-ip-address      4559      active
```

As described in the ASE document [Configuring the Server for Multiple Network Handlers](#), this configuration is read from the `$SYBASE/interfaces` file on your ASE server during instance startup.

Resolution

To allow discovery of instances using multiple network interfaces, use the **Add Environment** dialog to add the environment ONCE per network interface.

Each environment added in this way should have both of the following:


- A unique Host Address; and
- A unique Toolkit Path

The Delphix Engine expects that the files deployed to the Delphix Toolkit directory are persistent. Any failover of the ASE service to a different host must ensure that the Delphix Toolkit directory remains consistent between both hosts, with all files and permissions intact.

Deploying the Delphix Toolkit multiple times will result in increased disk space consumption. It may circumvent concurrency limitations that the Delphix Engine introduces to minimize its impact on your source environments. If your environment requires you to add the same host several times using this process, contact your Delphix

Professional Services representative or Customer Success Manager to discuss best practices or alternative solutions.

Using the example interfaces file provided above, you could use the following inputs to add the `INSTANCENAME` instance:



The screenshot shows the 'Add Environment' web interface. On the left, there is a navigation menu with three items: 'Host and Server', 'Environment Settings' (which is selected and highlighted with a blue dot), and 'Summary'. The main content area is titled 'Add Environment' and contains several form fields and sections:

- DSP TrustStore Password**: A text input field.
- Login Type**: Two radio button options: 'Username and Password' (selected) and 'Username and Public Key'.
- OS Username**: A text input field.
- OS Password**: A text input field.
- Validate**: A button.
- Toolkit Path**: A text input field.
- Discover SAP ASE**: A checked checkbox labeled 'Enabled'.
- ASE DB Username**: A text input field containing the text 'delphix_db'. This field is highlighted with a blue border.
- ASE DB Password**: A text input field with masked characters (dots).
- Notes**: A text input field.

Related /topics

- [Managing SAP ASE Environments and Hosts](#)

Changing the Host Name or IP Address of an SAP ASE Environment

This topic describes how to change the hostname or IP address for source and target environments, and for the Delphix Engine.

Procedure

For Source Environments

1. Disable the dSource as described in [Enabling and Disabling SAP ASE dSources](#).
2. If the **Host Address** field contains an IP address, edit the IP address.
3. If the **Host Address** field contains a hostname, update your Domain Name Server to associate the new IP address to the hostname. The Delphix Engine will automatically detect the change within a few minutes.
4. In the **Environments** screen of the Delphix Engine, refresh the host.
5. Enable the dSource.

For VDB Target Environments

1. Disable the VDB as described in [Enabling and Disabling SAP ASE dSources](#).
2. If the **Host Address** field contains an IP address, edit the IP address.
3. If the **Host Address** field contains a hostname, update your Domain Name Server to associate the new IP address to the hostname. The Delphix Engine will automatically detect the change within a few minutes.
4. In the **Environments** screen of the Delphix Engine, refresh the host.
5. Enable the VDB.

For the Delphix Engine

1. To stop running your VDB select the red **Stop** button located on the VDB **Configuration** tab.
2. Disable all dSources as described in [Enabling and Disabling SAP ASE dSources](#).
3. You can use either the command-line interface or the Delphix Setup application to change the IP address of the Delphix Engine.
 - a. To use the command-line interface, press **F2** and follow the instructions described in [Setting Up Network Access to the Delphix Engine](#).
 - b. To use the Delphix Setup application, go to **Delphix Management > Engine Setup** in the Delphix Management application, or click **Server Setup** in the Delphix Engine login screen.
 - i. In the **Network** panel, click **Modify**.
 - ii. Under **DNS Services**, enter the new IP address.
 - iii. Click **OK**.
4. Refresh all Environments by clicking the **Refresh** symbol on the **Environments** screen.
5. Enable all dSources as described in [Enabling and Disabling SAP ASE dSources](#).
6. To start all VDBs, click the **Start** button located on the VDB **Configuration** tab.

Related Topics

- [Managing SAP ASE Environments](#)

Using HostChecker to Validate SAP ASE Source and Target Environments

What is HostChecker?

The HostChecker is a standalone program which validates that host machines are configured correctly before the Delphix Engine uses them as data sources and provision targets.

Please note that HostChecker does not communicate changes made to hosts back to the Delphix Engine. If you reconfigure a host, you must refresh the host in the Delphix Engine in order for it to detect your changes.

You can run the tests contained in the HostChecker individually, or all at once. You must run these tests on both the source and target hosts to verify their configurations. As the tests run, you will either see validation messages that the test has completed successfully, or error messages directing you to make changes to the host configuration.

Prerequisites

- Make sure that your source and target environments meet the requirements specified in [SAP ASE Support and Requirements](#).

Procedure

1. Download the HostChecker tarball matching the host's operating system from <https://download.delphix.com/> (for example: `hostchecker_linux_x86.tar`).
2. Create a working directory and extract the HostChecker files from the HostChecker tarball.

```
mkdir dlp-host-checker
cd dlp-host-checker/
tar -xf hostchecker_linux_x86.tar
```

3. Change to the working directory and enter this command. Note that for the target environments, you would change `source` to `target`.

```
$ cd hostchecker
$ ./hostchecker.sh
Extracting the JDK from the tarball jdk.tar.gz.
Please enter whether this machine is a source or a target:target
1: Check ASE environment
2: Check all ASE instances
3: Check all the Oracle installations
4: Check homedir permissions
5: Check Linux Performance Settings
6: Check mkdir and rmdir
7: Check the MySQL installation
8: Check network port access
9: Check the Oracle CRS home
10: Check for ssh connectivity
11: Check sshd_config for timeout configuration
12: Check user sudo privileges
13: Check toolkit path
all: Execute all checks
quit: Exits
```

Please **select** an option:



Don't Run as Root

Don't run the HostChecker as root; this will cause misleading or incorrect results from many of the checks.

4. Select which checks you want to run. We recommend you run all checks (excluding Oracle and MySQL) if you are running Hostchecker for the first time.
5. Pass in the arguments the checks ask for.
6. Read the output of the check.
7. The error or warning messages will explain any possible problems and how to address them. Resolve the issues that the HostChecker describes. Don't be surprised or undo your work if more errors appear the next time you run HostChecker, because the error you just fixed may have been masking other problems.
8. Repeat steps 3–7 until all the checks return no errors or warnings.

Tests Run

Test	SAP ASE Source	SAP ASE Target	Description
Check SAP ASE environment	X	X	<ul style="list-style-type: none"> • Checks that the \$SYBASE environment variable exists for interactive logs. NOTE: It is still necessary to check that it is defined for non-interactive logins (ssh user@host env grep SYBASE). • Checks that the "isql_r64" binary can be found underneath the \$SYBASE environment variable.
Check All SAP ASE instances	X	X	<p>Attempts to connect to each running SAP ASE instance via "isql_r64" (utilizing the \$SYBASE/interfaces file) and execute the following queries:</p> <pre>select @@servername select count(*) from master..syslisteners select count(*) from master..sys.servers select count(*) from master..sys.databases select count(*) from master..sys.usages select @@version select @@maxpagesize SELECT srvnetname FROM master..sys.servers WHERE srvname='SYB_BACKUP'</pre>

Test	SAP ASE Source	SAP ASE Target	Description
Check all the Oracle installations	N/A	N/A	
Check homedir permissions	X	X	Check that the home directory, the ~/.ssh directory, and the ~/.ssh/authorized_keys file exist, that they are owned by the user invoking this check, and that they are not 'group' or 'other' writable.
Check <OS> Performance Settings		X	Check target's kernel settings necessary to optimize performance.
Check mkdir and rmdir		X	Tests that the user can mkdir and rmdir under both the toolkit directory and the specified mount path.
Check the MySQL installation	N/A	N/A	
Check network port access	X	X	Can be used to test access to specified ports on the Delphix Engine. See Network and Connectivity Requirements for SAP ASE Environments for a list of ports.
Check the Oracle CRS home	N/A	N/A	
Check Oracle DB Instance	N/A	N/A	
Check for ssh connectivity	X	X	Verifies that the environment is accessible via SSH
Check sshd_config for timeout configuration		X	Check that sshd_config does not contain ClientAliveInterval or ClientAliveCountMax entries.
Check toolkit Path	X	X	Verifies that the toolkit installation location has the proper ownership, proper permissions, and enough free space.

Test	SAP ASE Source	SAP ASE Target	Description
Check user sudo privileges	X	X	Verifies that the operating system user can execute certain commands with necessary privileges via <code>sudo</code> . This only needs to be run on target environments. See the topic Requirements for SAP ASE Environments and Databases for more information.

Related Topics

- [SAP ASE Support and Requirements](#)
- [Managing SAP ASE Environments](#)
- [Requirements for SAP ASE Environments and Databases](#)

Enabling Linking and Provisioning for SAP ASE Environments

This topic describes how to enable and disable provisioning and linking for SAP ASE databases.

Before a database can be used as a dSource, you must first make sure that you have enabled linking to it. Similarly, before you can provision a VDB to a target database, you must make sure that you have enabled provisioning to it.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **Databases** tab.
5. Click the **pencil** icon located next to the database **Installation Details**.
6. Select the **Allow Provisioning** checkbox to enable provisioning, deselect the checkbox to disable provisioning.
7. Click **show details** for the database.
8. Slide the button next to **Allow Linking** to **On** or **Off** to enable or disable linking.

Related Topics

- [Managing SAP ASE Environments Overview](#)
- [SAP ASE Environments and Data Sources](#)

Linking Data Sources and Syncing Data with SAP ASE

Creating a dSource will ingest data from the source and create a dSource on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

This section covers the following topics:

- [Linking SAP ASE Data Sources: An Overview](#)
- [Linking Data Sources with SAP ASE](#)
- [Linking Data Sources with Encrypted SAP ASE Database](#)
- [Data Management Settings for SAP ASE dSources](#)
- [Upgrading a SAP ASE dSource](#)
- [Detaching and Re-attaching SAP ASE dSources](#)
- [Deleting an SAP ASE dSource](#)
- [Enabling and Disabling SAP ASE dSources](#)
- [Working with SAP ASE Snapshots](#)

Linking SAP ASE Data Sources: An Overview

This topic describes basic concepts behind the creation of dSources from SAP ASE databases.

Initial Linking and Staging Databases

A dSource is a copy of a physical database that is created when the Delphix Engine links to and loads the database from a backup. The database backup can be a new full database backup that the Delphix Engine initiates, the most recent existing database backup, or an existing database backup specified by the user. When loading from an existing backup, the backup should be in a location that the source environment user can access.

After obtaining the initial snapshot and linking the dSource, the Delphix Engine keeps the dSource and the source database in sync by monitoring the source database for new transaction log dumps and then applying those backups on a standby database. This database is called the "staging database." A target environment that hosts one or more staging databases is referred to as a "staging target."

After you have linked a database into the Delphix Engine, you can re-initialize it by performing a sync on the dSource.

 There is an SAP ASE feature that affects the size of database dumps which correspondingly affect the amount of space consumed on the storage attached to the Delphix Engine. If the SAP ASE "sp_dumpoptimize" feature is set to maximum as follows, more disk space will be consumed on Delphix storage because SAP ASE writes both allocated and unallocated pages into the dump files:

```
sp_dumpoptimize 'archivespace = maximum'
```

SAP ASE Source and Staging Database Compatibility

You may use different SAP ASE patch level versions between source and staging hosts. The major version of SAP ASE must still be the same. The only caveat is for SAP ASE 15.7 - where the source and staging patch levels must both be either any version below SP64 or both be any version above SP100.

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Linking Data Sources with SAP ASE](#)
- [Adding an SAP ASE Environment](#)

Linking Data Sources with SAP ASE

The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools. For an overview of all dSource related actions, please visit [Managing Data Sources](#).

Delphix Virtualization for SAP ASE databases leverages backup-based ingestion, which means that Delphix will look for, or sometimes initiate the creation of, a backup through your SAP ASE backup server. From there, the backup is restored on a staging server and the staging copy is then ingested into Delphix. See [Delphix Architecture with SAP ASE](#) for more information.

When linking a dSource from an SAP ASE source database, Delphix offers several different methods of capturing backup information:

- ASE Managed Backups, where the SAP ASE source database schedules and initiates backups. This method supports various backup types which include:
 - Full backups
 - Transaction log backups (with LogSync disabled)
 - Transaction log backups (with LogSync enabled)
- Delphix Managed Backups, where the Delphix Engine schedules and initiates the backups from the source database, and captures them.

ASE Managed Backups

Further contextual information on the various backup types (listed above):

- Full Backups - A snapshot will be created on the Delphix Timeflow for each Full backup.
- Transaction log backups (with LogSync disabled) - A snapshot will be created on the Delphix Timeflow for each transaction log backup.

Transaction Logs

Transaction logs are not collected if:

- a) there are gaps in the sequence of log backups (a break in the “log chain”).
- b) the available log backups do not include any changes since the last successful Delphix snapshot.

- Transaction log backups (with LogSync enabled) - A snapshot will be created on the Delphix Timeflow for each transaction log backup. In addition, point-in-time provisioning will be an available option if you would like to provision from any point in between snapshots.

Log Files

Log files consume additional space on the Delphix Engine and are managed according to the defined retention policy for logs.

Delphix Managed Backups

When the checkbox for Delphix Managed Backups is selected, the Delphix Engine will initiate a full backup of the source database for the initial load of the dSource. Thereafter, the Delphix Engine will initiate full backups of the source database using the schedule specified by the selected SnapSync Policy. If you select the None policy, the Delphix Engine will not automatically initiate a full backup, but you can initiate them manually using the snapshot (camera) icon.

Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment)

Procedure

1. Login to the **Delphix Management** application.
2. Navigate to **Manage > Datasets**.
3. Click the plus icon and select **Add dSource**.

 Delphix looks for the backup files on the staging host (unless a "remote" backup server is used in which case, the remote host is used which is often the source environment).

4. In the **Add dSource** wizard, select the source environment with the correct environment-based user.
5. Enter your login credentials for the source database and click **Next**.
6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Select the **Data Management** settings needed, as described in [Data Management Settings for ASE Data Sources](#).
8. From the **Data Management** tab under the **Initial Load** option, select and enter any **additional settings** needed. There are three options for the initial load of the dSource:
 - a. If the source ASE instance resides on the same server as the staging ASE instance, the staging database's NFS mounted "temp" directory will be present for the source database to write to in response to the "DUMP DATABASE" command.
 - b. If the source and staging ASE instances are configured to allow remote access to the backup servers and the dSource is linked using the **Remote Server** option as described below. Delphix will then issue the "DUMP DATABASE" command and append the "AT <staging_backup_server_name>" clause so that the dump is written to the staging backup server.
1. (Recommended) **New Full Backup** - Lets Delphix create a new full backup file and load it. Note - that when Delphix creates the backup, it is moved to Delphix's NFS-mounted storage located on the stage host rather than. The backup will be located in the "temp" directory and will be deleted once the Delphix Engine has restored the backup and created a dSource from the restored staging database. This means that this option will work under two scenarios:
2. Most Recent **Existing Full Backup** – Find the most recent existing full backup file in the Backup Location and load it.

 If Dump History is not active on the Source Database: Choosing this option can delay completion of the dSource link as Delphix attempts to find and catalog every single backup listed in the source database's backup server log file.

3. Specific **Existing Full Backup** – Specify which backup files in the Backup Location you want to load. Choosing this option is much faster because Delphix will skip directly to loading the desired backup and only start to search for and catalog backups in the background after the linking of the source database has completed.

 When using a dump taken with the deprecated compression syntax, select the Specific Existing Full Backup option for Initial Load and, for each stripe, type compress::<file name> into the text box.

- Select the Staging environment and ASE instance that will be used to manage the staging database used for validated sync of the dSource.
- Select any policies for the new dSource.
- Click **Next**, then specify any pre-hook and post-hook scripts.

- Review the dSource Configuration and Data Management information, and click **Submit** to begin provisioning the VDB.

Related Topics

- [Data Management Settings for SAP ASE dSources](#)
- [Managing Data Sources and Syncing Data](#)

Linking Data Sources with Encrypted SAP ASE Database

This topic describes the behavior of the Delphix Engine when linking to a dSource based on an encrypted SAP ASE database.

Beginning with version 16.0, SAP ASE supports the Database Encryption feature. You can link the encrypted source databases to a dSource by following the basic procedure described in [Linking Data Sources with SAP ASE](#) but the user needs to ensure that the same key with the same name is present on the staging instance as it was on the source instance. The keys should be imported from the source instance. Refer to the configuration steps 4 and 5 of [Delphix Implementation of Database Encryption](#) to export and import the master and encryption key.



- If a user links a database when it is unencrypted and then encrypts it, the subsequent syncs will fail.
- If a user links an encrypted database and then decrypts the source database, the subsequent syncs will fail.
- If a user changes the encryption key, the subsequent syncs will fail.

In the case of the above failures, the user will be prompted to apply a full backup with the “Drop and Recreate Devices” option enabled (for once), and subsequent snap and validate sync will run without any failures. See [Working with Snapshots](#).



Exceptions with suggested actions are raised in the following scenarios:

- If there has been a change in the encryption status of a source database since it was last synced or the key has changed.
- If the required encryption key is not present on the staging instance during the linking.

Data Management Settings for SAP ASE dSources

Each dSource has its own data management settings, which can be configured during the linking workflow as well as on the configuration page for that dSource. You can configure data management settings to improve overall performance and meet your requirements.

The following settings are available for SAP ASE data sources:

Setting	Explanation
Backup Path	The directory where Delphix will search for backups for data ingestion. Delphix recursively searches this location, so the database backups or transaction logs can reside in any subdirectories below the path entered.
Validated Sync Mode	<p>ASE Validated Sync is either enabled or disabled. Delphix stays in sync with source databases by monitoring the ASE backup server log or dump history file (depending on the dSource configuration). When it sees a new database dump or transaction log has been created, it attempts to load it into the ASE staging instance.</p> <ul style="list-style-type: none"> • Enabled - Delphix automatically ingests either transaction logs or full database dumps depending on the source database's configuration. • Disabled - Delphix does not actively monitor the source database's backup server log or dump history file and does not automatically ingest backups.
Dump History	Delphix Validated Sync will leverage Dump History files to locate and ingest backups. This setting is recommended to improve the data ingestion and syncing processes.
LogSync	<p>When enabled, Delphix copies the transaction logs to Delphix storage to allow precise point-in-time provisioning.</p> <p>Note: If Dump History is not enabled, due to SAP ASE CR 800569, Delphix can only support transaction logs generated in intervals greater than one minute apart for SAP ASE versions 16.0 SP02 through SAP ASE 16.0 SP02 PL04. This ASE bug inadvertently removed the second and millisecond precision from the dump header sequence dates preventing Delphix from knowing what order to apply the transaction logs when there are multiple transaction logs dumped within the same minute.</p>
Source of Production Backup	<p>You may select a staging or remote server as the location for the Delphix Engine to ingest backups.</p> <p>Note - Selecting a staging server is the most common application for ingesting source-based backups. One major reason for this is the ability to use Validated Sync to coordinate backup ingestion.</p>

Setting	Explanation
Dump Password	SAP ASE has an optional field for database dumps (backups) to prevent unauthorized loads. If enabled, Delphix will need to store the dump password in order to restore the SAP ASE backup for ingestion. Delphix also provides the option to propagate password protection to established dSources in the case of an active password on the originating database dump or transaction log dump files. Setting this option causes Delphix to add the "WITH passwd=" clause to the "LOAD" commands.

Procedure

1. Enter the **Backup Location**. This is the directory where the database backups are stored. Delphix recursively searches this location, so the database backups or transaction logs can reside in any subdirectories below the path entered.
2. Select the **Staging Environment** and ASE instance name.
3. Enable or disable Validated Sync Mode. Validated Sync Mode (also known as ValidatedSync) is a background process that looks for new full or transaction log backups either by monitoring ASE Backup Server's log file or Dump History file for a new database or transaction log dumps. When Delphix detects a new dump is available it attempts to load it into the staging ASE database.
4. Enable or disable **Use Dump History**. If Dump History is enabled, Delphix uses the ASE Dump History file to find information about backups being performed on the source database and to find the next backup to be used for Validated Sync.
5. Enable or disable **LogSync**. LogSync copies the transaction logs to Delphix storage which enables provisioning VDBs from a specific point-in-time in rather than just a particular backup

LogSync support limitations

[If Dump History is not active on the Source Database]

Due to ASE CR 800569, Delphix can only support transaction logs generated in intervals greater than one minute apart in ASE versions 16.0 SP02 through ASE 16.0 SP02 PL04. This ASE bug inadvertently removed the second and millisecond precision from the dump header sequence dates preventing Delphix from knowing what order to apply the transaction logs in when there are multiple transaction logs dumped within the same minute.

6. Select **Backup Location Type**.
7. Click **Advanced** to edit Source of Production Dump, External Data Directory, Retention policies, or Dump Password.

External Data Directory

The External Data Directory feature is not currently used with ASE dSources and is targeted for removal in a future release of Delphix.

Remote Server should be selected when database dumps cannot be found on the Staging Environment. This option can be used with any of the initial load selections (New Full Backup, Most Recent Existing Full Backup or Specific Existing Full Backup). If selected, fill out additional settings as needed:

- a. Enter the Remote Server Name. This is the name of the backup server used when the dump was created.
- b. Select the Remote Host and Remote User that the backup server is located on.
- c. As noted, the interfaces file on both the staging and remote environments must be modified to point at each other's backup servers.

- d. The Create Dump Password sets a dump password for the dSource. Select this only if the dump password option was used to create a password on the database dump or transaction log dump files. Setting this option causes Delphix to add the "WITH passwd=" clause to the "LOAD" commands.

Related Topics

- [SAP ASE Environments and Data Sources](#)
- [Linking Data Sources and Syncing Data with SAP ASE](#)

Upgrading a SAP ASE dSource

This topic describes how to upgrade SAP ASE dSources.

Prerequisites

- The staging SAP ASE instance must be the same version as the SAP ASE instance hosting the upgraded source database.

Procedure

1. Login to the **Delphix Management** application.
2. In the top menu bar, click **Manage** and select **Datasets**.
3. Select the **dSource** to be upgraded.
4. Disable the dSource.
 - a. From the **Actions (...)** menu select **Disable**.
 - b. In the Disable dialog select **Disable**.
5. Upgrade the source and staging ASE instances (Delphix requires the source and staging ASE instances to be at the same major release level).
6. Click the **Refresh** icon (under the **Manage > Environments** menu) to refresh the Delphix environment hosting the source and staging ASE instances so that Delphix registers the new version of ASE.
7. Enable the dSource.
 - a. From the **Actions (...)** menu select **Enable**.
 - b. In the Enable dialog select **Enable**.

Additional steps

If Delphix was unable to discover the SAP ASE instances automatically using Manual Discovery, you will need to manually update the version by completing the following procedure, [CLI Cookbook: Configuring SAP ASE Manual Discovery](#).

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Upgrading SAP ASE VDBs](#)

Detaching and Re-attaching SAP ASE dSources

This topic describes how to detach dSources and re-attach them to a different source database.

Each dSource contains metadata that associates it with the source database, as well as the data it has ingested from the source database in the form of snapshots up to that point. It is possible to detach, or unlink, a dSource from its source database. This breaks the association with the source database without affecting the data within the Delphix Engine. Detached dSources and their source databases have these properties:

- A detached dSources can still be used to provision a virtual database (VDB).
- You can re-link the source database as a different dSource.
- Any child VDBs that were provisioned from this dSource will only be able to be refreshed from the most recent snapshot available on the dSource.
- If you need a VDB from a newer snapshot, you would need to provision a new VDB. Once you have provisioned the new VDB you can delete the old VDBs provisioned from this dSource. You can delete the old dSource when it is no longer needed.

Detaching a dSource

1. Login to the **Delphix Management** application as a user with **OWNER** privileges on the dSource, group, or domain.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **Dataset** you want to unlink.
5. From the **Actions** menu (...) select **Unlink dSource**.
A warning message will appear.
6. Click **Unlink** to confirm. The status of the dSource will show as Detached.

Rebuilding Source Databases and Using VDBs

In situations where you want to rebuild a source database but retain the existing dSource, you will need to detach the original dSource and create a new one from the rebuilt data source.

1. Detach the dSource as described in the procedure on this page.
2. You cannot attach a dSource with the same name as a dSource that is already attached. If you intend to give the new dSource the same name as the original one, rename the detached dSource.
 - a. At the top of the **Configuration** tab, next to the dSource's name, click the **Edit** (pencil) icon.
 - b. After renaming the dSource, click the **checkmark**.
3. Create the new dSource from the rebuilt database.

You will now be able to provision VDBs from both the detached dSource and the newly created one, but the detached dSource will only represent the state of the source database prior to being detached.

Attaching a Previously Detached dSource

The attach operation is supported via the UI. You can only re-attach databases that represent the same physical database.

Attaching a dsouce via the UI:

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **Dataset** you want to link.
5. From the **Actions** menu (...) select **Link dSource**.
6. In the Link dSource window click **Link** to confirm.

Link dSource ✕

Source Environment
bbdhcp-LP29-79440.dcenter.delphix.com

Installation
/u01/app/ora10205/product/10.2.0/db_1

Database
dbdhcp1

Environment User
sybase

Database Username
admin

Database Password

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)

Deleting an SAP ASE dSource

This topic describes how to delete an SAP ASE dSource.

Deleting a dSource will delete the dSource metadata for a particular source database, along with all snapshots, logs, and policies stored in Delphix. This is a permanent operation, and re-attaching the dSource will require [a new linking operation](#). If a dSource is deleted, it does not affect your source database.

If you wish to temporarily disable your dSource without deleting it, you can follow the steps to [Enabling and Disabling SAP ASE dSources](#) instead.

Prerequisites

You cannot delete a dSource that has dependent virtual databases (VDBs). Before deleting a dSource, make sure that you have deleted all dependent VDBs as described in [Deleting a VDB](#).

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Click **Datasets**.
4. In the **Datasets** list, select the **dSource** you want to delete.
5. From the **Actions** menu (...) select **Delete**.
6. Click **Delete** to confirm.

 Deleting a dSource will also delete all snapshots, logs, and descendant VDB refresh policies for that dSource. You cannot undo the deletion.

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)
- [Detaching and Re-attaching SAP ASE dSources](#)
- [Deleting an SAP ASE VDB](#)

Enabling and Disabling SAP ASE dSources

This topic describes how to enable and disable dSources when certain operations against the source database must occur outside of Delphix.

Some operations, such as restoring the source database from a backup, will require that the dSource be temporarily disabled. Disabling a dSource turns off communication between it and the source database, but it does not tear down the configuration that enables communication and data updating to take place. When a disabled dSource is later enabled, it will resume communication and incremental data updates from the source database according to the original policies and data management configurations that you set.

Disabling a dSource is also a prerequisite for several other operations, such as database migration and upgrading the dSource metadata after upgrade of the associated data source.

Procedure

Disabling a dSource will stop further operations on the Delphix Engine related to the dSource.

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **dSource** you want to disable.
5. In the upper right-hand corner, from the **Actions** menu (...) select **Disable**.
6. In the Disable dialog select **Disable**.

When you are ready to enable the dSource again, from the Actions menu (...) select **Enable**, and the dSource will continue to function as it did previously.

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)

Working with SAP ASE Snapshots

Taking a snapshot creates a new snapshot entry in the dSource's Timeflow.

Snapshot

This section lists the steps to take a snapshot and delete the same.

Login to the **Delphix Management** application.

1. Click **Manage** and select **Datasets** from the dropdown list.
2. Select the dSource you want to Snapshot.
3. Click the **Camera** icon.
4. Select the method to take the snapshot. See [Quick Start Guide for SAP ASE](#) for information on Delphix Managed Backups.

You can optionally select the **Enable** checkbox under **Drop and recreate devices**. This option can be used if the previous sync failed because of device mapping issues. When the sync operation is performed with this option set to true, Delphix Engine will not try to re-map the existing devices; instead, it will delete them and create new devices on the staging database. This option should only be used if the previous sync had failed due to device remapping, as this will increase disk usage.

5. From the Snapshot dialog, click **Snapshot**.
6. Under the **Timeflow** tab, choose an option to view the snapshot and verify the snapshot you just created.
7. To delete, select a snapshot you want to delete, and from the Actions menu (...) select **Delete Snapshot**.
8. Under the **Timeflow** tab, choose an option to view the snapshot and verify the snapshot you just deleted.

Provisioning and Managing VDBs from SAP ASE

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment.

From a dSource, you can select a snapshot or point in time to create a VDB. SQL Server VDBs each have their own configuration settings as described in Configuration Settings for ASE Virtual Databases below.

This section covers the following topics:

- [Overview of Provisioning SAP ASE Virtual Databases](#)
- [Provisioning an SAP ASE VDB](#)
- [Provisioning a VDB from an Encrypted SAP ASE Database](#)
- [Provisioning an SAP ASE VDB from a Replicated VDB or dSource](#)
- [Resizing an SAP ASE VDB](#)
- [V2P with an SAP ASE dSource or VDB](#)
- [Additional SAP ASE VDB Topics](#)

Overview of Provisioning SAP ASE Virtual Databases

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment, as described in the overview for [Managing Environments and Hosts](#) and [Requirements for SAP ASE Environments and Databases](#).

From a dSource, you can select a snapshot or point in time to create a VDB. SQL Server VDBs each have their own configuration settings as described in Configuration Settings for ASE Virtual Databases below.

Procedure

1. In the Datasets panel on the left-hand side, click the group containing the dSource or VDB from which you want to provision.
2. Click the **Timeflow** tab.
3. Select a snapshot or open LogSync timeline to provision by a specific log or point in time.
 - Find more detail about initial provisioning options in the section ‘Provisioning by Snapshot or LogSync’ below.
 - Click to open the Provision VDB wizard, and select a compatible Target Environment for the new ASE VDB
 - Review the information presented for Target Configuration and edit as necessary.
 - Select a Snapshot Policy for the VDB.
 - (Optional) - Selective Data Distribution - After policies, there is a masking option.
 - Enter any operations that should be run in the Hooks page. These scripts can be managed after provision in the VDB’s configuration page.

When provisioning starts, you can review the progress of the job by selecting the VDB and clicking on the Status tab, or by selecting Manage/Dashboards and viewing the Job History panel. Alternatively, you could see this in the Actions Sidebar. When provisioning is complete, the VDB will be included in the group you designated and listed in the Datasets panel. If you select the VDB in the Datasets panel and click the Configuration tab, you can view information about the database and its Data Management settings.

Provisioning by Snapshot or LogSync

When provisioning by Snapshot, you can provision to the start of any particular snapshot by time.

Provisioning By Snapshot/Time	Description
Provisioning By Snapshot	You can provision by using a Snapshot. In that case, a new VDB will be provisioned to the database state as of the Snapshot.
Provision by Time	If you have enabled Log Sync, you can provision a new database to a point in time. You can select a snapshot and then using time entry fields, specify a Point in Time. Delphix will use the selected snapshot to restore the VDB and use the log files to roll forward the VDB to the selected time.

Configuration Settings for ASE Virtual Databases

Each VDB has its own data management settings, found during the provisioning workflow as well as in the configuration page for that VDB. When you create a SAP ASE VDB, Delphix copies most configuration settings from the dSource and uses them to create the VDB. However, you can customize these with the following settings:

Setting	Explanation
Recovery Model	The current recovery model of the source database. This field will auto-populate with information from the dSource.
Auto VDB Restart	Enabling this option will automatically restart this VDB whenever its target host is rebooted.

Automatic VDB Restart on Target Server After Reboot

The Delphix platform now automatically detects whether a target server has been rebooted, and proactively restarts any VDB on that server that was previously up and running. This is independent of the data platform. It is done as if you realized a target server was restarted and issued a start command from the Delphix platform. This feature is compatible with Self-Service ordering dependencies and is limited to non-clustered VDBs.

To enable automatic restart, complete the following steps:

- When provisioning a new VDB in the VDB Provisioning wizard, check the **Auto VDB Restart** box.

Once the VDB has been provisioned, you will be able to turn **Automatic VDB Restart** on.

1. In the **Datasets** panel, select the VDB.
2. Select the **Configuration** tab.
3. Select **Source** sub-tab.
4. Select **Database edit**.

Related Topics

- [Provisioning and Managing Virtual Databases](#)
- [Provisioning and Managing VDBs from SAP ASE](#)

Provisioning an SAP ASE VDB

This topic describes how to provision a virtual database (VDB) from an SAP ASE dSource.

Prerequisites

- You must have already linked a dSource from a source database, as described in [Linking Data Sources with SAP ASE](#), or have already created a VDB from which you want to provision another VDB.
- You must have already set up target environments as described in [Adding an SAP ASE Environment](#).
- Ensure that you have the required privileges on the target environment, as described in [Requirements for ASE Environments and Databases](#).
- If you are provisioning to a target environment that is different from the one in which you set up the staging database, you must make sure that the two environments have compatible operating systems, as described in [Requirements for SAP ASE Target Hosts and Databases](#). For more information on the staging database and the validated sync process, see [Managing SAP ASE Environments Overview](#).

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**
3. Select **Datasets**.
4. Select a **dSource**.
5. Select a **means of provisioning**.
6. Click **Provision**.
The **Provision VDB** panel will open, and the **Instance** and **Database Name** fields will auto-populate with information from the dSource.
7. Select whether to enable **Truncate Log on Checkpoint** database option for the VDB.
8. Click **Next**.
9. Select a **Target Group** for the VDB.
Click the green **Plus** icon to add a new group, if necessary.
10. Select a **Snapshot Policy** for the VDB.
Click the green **Plus** icon to create a new policy, if necessary.
11. Click Auto VDB Restart to enable VDBs to be automatically restarted when staging/target host gets rebooted, if necessary.
12. Specify any **Hooks** to be used during the provisioning process.
For more information, see [SAP ASE Hook Operations](#).
13. If your Delphix Engine system administrator has configured the Delphix Engine to communicate with an SMTP server, you will be able to specify one or more people to notify when the provisioning is done. You can choose other Delphix Engine users or enter email addresses.
14. Click **Submit**.
When provisioning starts, the VDB will appear in the **Datasets** panel. Select the VDB and navigate to the **Status** tab to see the progress of the job. When provisioning is complete, more information can be seen on the **Configuration** tab.

Related Topics

- [Linking Data Sources with SAP ASE](#)
- [Adding an SAP ASE Environment](#)
- [Requirements for SAP ASE Environments and Databases](#)

Provisioning a VDB from an Encrypted SAP ASE Database

This topic describes how to provision a VDB from an encrypted database.

The Delphix Engine supports provisioning from a dSource linked to a physical database that has been encrypted with SAP ASE Database Encryption, which can be used to encrypt columns or complete databases.

You can provision a VDB from a dSource linked to an encrypted database using the normal process to [Provision an SAP ASE VDB](#) but the user needs to ensure that the same key with the same name is present on the target instance as it was on the staging instance because we mount the database onto the target that was created on the staging instance. The keys should be imported from the source instance. Refer to configuration steps 4 and 5 of [Delphix Implementation of Database Encryption](#) to export and import the master and encryption key.

-  Exception with suggested actions is raised in the following scenario:
- If the required encryption key is not present on the target instance during the provisioning.

Provisioning an SAP ASE VDB from a Replicated VDB or dSource

This topic describes how to provision from a replicated dSource or virtual database (VDB). The process for provisioning from replicated objects is the same as the typical VDB provisioning process, except that first you need to select the **replica** containing the replicated object.

Prerequisites

- You must have replicated a dSource or a VDB to the target host, as described in [Replication Overview](#)
- You must have added a compatible target environment on the target host

Procedure

1. Login to the **Delphix Management** application for the target host.
2. Click **Manage**.
3. Select **Datasets**.
4. In the list of replicas, select the **replica** that contains the dSource or VDB you want to provision.
5. The provisioning process is now identical to the process for provisioning standard objects.

Post-Requisites

Once the provisioning job has started, the user interface will automatically display the new VDB in the live system.

Related Topics

- [Provisioning and Managing VDBs from SAP ASE](#)
- [Replication Overview](#)

Resizing an SAP ASE VDB

When resizing ASE VDBs, we recommend:

- Ensure that the VDB's database devices reside on the Delphix storage, in the same directory as the database's other devices. Use the ASE "sp_helpdb" and "sp_helpdevice" commands to get details on the VDB's devices.
- If you need to add a new device, use the "skip_alloc=true" clause because the storage on Delphix is already zeroed (initialized). This should allow the "DISK INIT" command to complete almost instantaneously.
- If you add a new device to the database by issuing the ASE "DISK INIT" followed by the "ALTER DATABASE" commands, click the camera icon to take a new snapshot of the VDB.

 **Note**

It is important to take a snapshot after altering the VDB's device layout so that Delphix rewrites the database's "manifest" file to include the new database device.

- Delphix 6.x uses the new "ALLOW_DBID_MISMATCH, FIXDBID" clause to fix in dbid mismatches on versions of ASE that support these clauses. If you are on an older version of Delphix or ASE that does not support this clause, you may be required to run the DBCC CHECKALLOC() command on the VDB.

If you use the ASE **ALTER DATABASE** command to increase the amount of space allocated to a VDB, you may need to run the **DBCC CHECKALLOC()** command on the VDB. If the VDB is unmounted and remounted without this command being run on it, you may run into the following ASE CR which leaves the VDB in an unusable state:

SAP ASE Change Request Number	Description
798271	<p>The errors 14545, 14547 and 14519 will be raised by MOUNT DATABASE if the unmounted database had previously been mounted on a server where the database ID was used and a new database ID was assigned, and this database was extended without previously running DBCC CHECKALLOC(dbname, fix). Additionally, two new options have been added to the MOUNT command: WITH FIXDBID, to instruct the MOUNT command to fix any possible database ID mismatch, and WITH ALLOW_DBID_MISMATCH, to prevent that MOUNT DATABASE fails if the database has different database ID values in the allocation pages.</p> <p>The fix for this issue has been ported to 15.7 SP138 and 16.0 SP02 PL05.</p>

Determining the Need to Run DBCC CHECKALLOC

DBID mismatches are most likely to occur under the following circumstances:

- When provisioning VDBs to the same ASE instance hosting the staging database of the dSource the VDB was provisioned from.
- When provisioning multiple copies of the same VDB to the same ASE instance.

Delphix has found the following technique as a reliable way to determine whether or not you need to run DBCC CHECKALLOC() after increasing the amount of space allocated to the VDB.

In the following example, a VDB has been provisioned back to the same ASE instance where the staging database was being hosted. The staging database had a **dbid=11865**, so ASE automatically assigns the VDB the next lowest available dbid (**dbid=22** in this case).

1. Resize the database:

```
1> SELECT name, dbid FROM sysdatabases WHERE name = 'Vdb5'
2> go
name                dbid
-----
Vdb5                 22

1> DISK RESIZE name="AAA4$ff367xUvd67P7II_db1_dev26", size="4M"
2> go

1> ALTER DATABASE Vdb5 ON AAA4$ff367xUvd67P7II_db1_dev26="4M"
2> go
Extending database by 1024 pages (4.0 megabytes) on disk
AAA4$ff367xUvd67P7II_db1_dev26
Database Vdb5 which is currently offline has been altered from size 2816 logical
pages (2816 physical pages) to 3840 logical pages
```

2. Run the following queries to determine if there is a dbid mismatch on the pages of the VDB. **DBCC PAGE()** requires the **sybase_ts_role** role.

```
1> -- Run a query to generate DBCC commands:
1> SELECT 'dbcc page (Vdb5,' + convert(varchar(5), lstart) + ', 0, 0)' FROM
master..sysusages WHERE dbid=db_id('Vdb5') AND segmap != 0
2> go

-----
dbcc page (Vdb5,0, 0, 0)
dbcc page (Vdb5,1536, 0, 0)
dbcc page (Vdb5,1792, 0, 0)
dbcc page (Vdb5,2816, 0, 0)

1> -- Turn on DBCC traceflag 3604 to direct output to stdout
2> dbcc traceon(3604)
3> GO

1> dbcc page (Vdb5,0, 0, 0)
2> dbcc page (Vdb5,1536, 0, 0)
3> dbcc page (Vdb5,1792, 0, 0)
4> dbcc page (Vdb5,2816, 0, 0)
5> go | grep dbid
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
```

```

pageno=0 dealloc_count=102 allocnode=0 ptnid=99 allocation_page dbid=11865 timestamp=
0000 00002730, segmap=0x00000003 (0x00000002
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=1536 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=11865 timestamp=
0000 00000001, segmap=0x00000004 (0x00000004
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=1792 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=11865 timestamp=
0000 00000001, segmap=0x00000003 (0x00000002
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=2816 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=22 timestamp=00
00 00000001, segmap=0x00000003 (0x00000002

```

Notice that three pages (0, 1536 and 1792) have the staging database's dbid (11865) and one page (2816) has the correct dbid (22) that matches the dbid of 22 stored in sysdatabases.

If all the dbid's match the dbid in sysdatabases, there is no need to run DBCC CHECKALLOC().

This means that DBCC CHECKALLOC() must be run on the VDB to correct the dbid mismatch:

```

1> dbcc checkalloc(Vdb5)
2> go
...etc...
11 allocation pages have been corrected to match database ID 22.

```

Note: In the code block above 11 pages were fixed.

Now if you run the **DBCC PAGE()** query again you can see that all of the dbid's match:

```

1> dbcc page (Vdb5,0, 0, 0)
2> dbcc page (Vdb5,1536, 0, 0)
3> dbcc page (Vdb5,1792, 0, 0)
4> dbcc page (Vdb5,2816, 0, 0)
5> go | grep dbid
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=0 dealloc_count=102 allocnode=0 ptnid=99 allocation_page dbid=22 timestamp=000
0 00002730, segmap=0x00000003 (0x00000002
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=1536 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=22 timestamp=00
00 00000001, segmap=0x00000004 (0x00000004
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=1792 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=22 timestamp=00
00 00000001, segmap=0x00000003 (0x00000002
    bmass_next=0x0 bmass_prev=0x0 bdbid=22
pageno=2816 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=22 timestamp=00
00 00000001, segmap=0x00000003 (0x00000002

```

Recovering from DBID Mismatch Errors

If a VDB containing mismatched dbids gets unmounted, it will likely get ASE errors 14545, 14547 and 14519 during the next attempt to **MOUNT** it and it will now be in an unusable state. Options for recovery include:

- For ASE versions that do not have the fix for ASE CR 798271, you may rewind the VDB to a snapshot taken prior to resizing the VDB. Any data changes made since this point in time will be lost.

On newer versions of ASE where ASE CR 798271 has been fixed, you can manually mount the VDB using the new **MOUNT** command options:

```

1> -- Get the device listing for the VDB from the manifest
2> MOUNT DATABASE Vdb5 FROM '/export/home/sybase/toolkit/564dfe60-40b6-
aec9-26fb-3baeb4e7b23b-vdb-19/datafile/manifest' WITH LISTONLY
3> go
[database]
Vdb5
[device]
'/export/home/sybase/toolkit/564dfe60-40b6-aec9-26fb-3baeb4e7b23b-vdb-19/datafile/
dxnff367xUvd67P7II_db1_dev26' =
'AAA4$ff367xUvd67P7II_db1_dev26'
'/export/home/sybase/toolkit/564dfe60-40b6-aec9-26fb-3baeb4e7b23b-vdb-19/datafile/
dx9279VtfPAvf0m5xY_db1_dev27' =
'AAA5$279VtfPAvf0m5xY_db1_dev27'

1> -- Mount the VDB using the WITH ALLOW_DBID_MISMATCH clause
2> MOUNT DATABASE [Vdb5] AS [Vdb5] FROM '/export/home/sybase/toolkit/564dfe60-40b6-
aec9-26fb-3baeb4e7b23b-vdb-19/datafile/manifest'
3> WITH ALLOW_DBID_MISMATCH USING
4> '/export/home/sybase/toolkit/564dfe60-40b6-aec9-26fb-3baeb4e7b23b-vdb-19/datafile/
dxnff367xUvd67P7II_db1_dev26' = 'AAA4$ff367xUvd67P7II_db1_dev26',
5> '/export/home/sybase/toolkit/564dfe60-40b6-aec9-26fb-3baeb4e7b23b-vdb-19/datafile/
dx9279VtfPAvf0m5xY_db1_dev27' = 'AAA5$279VtfPAvf0m5xY_db1_dev27'
6> go
The physical device '/export/home/sybase/toolkit/564dfe60-40b6-
aec9-26fb-3baeb4e7b23b-vdb-19/datafile/dxnff367xUvd67P7II_db1_dev26'
has been automatically assigned the logical device name
'AAA6$ff367xUvd67P7II_db1_dev26'.
The physical device '/export/home/sybase/toolkit/564dfe60-40b6-
aec9-26fb-3baeb4e7b23b-vdb-19/datafile/dx9279VtfPAvf0m5xY_db1_dev27'
has been automatically assigned the logical device name
'AAA7$279VtfPAvf0m5xY_db1_dev27'.
Started estimating recovery log boundaries for database 'Vdb5'.
Database 'Vdb5', checkpoint=(1543, 13), first=(1543, 13), last=(1543, 13).
Completed estimating recovery log boundaries for database 'Vdb5'.
Started ANALYSIS pass for database 'Vdb5'.
Completed ANALYSIS pass for database 'Vdb5'.
Started REDO pass for database 'Vdb5'. The total number of log records to process is
1.
Completed REDO pass for database 'Vdb5'.
MOUNT DATABASE: Completed recovery of mounted database 'Vdb5'.
MOUNT DATABASE: A new database id was required for database 'Vdb5' in order to mount
it. Execute DBCC CHECKALLOC(Vdb5, fixdbid) to
correct it.

1> DBCC CHECKALLOC(Vdb5, fixdbid)
2> go
Checking Vdb5: Logical pagesize is 4096 bytes

```

```
Total (# alloc pages = 15, # of alloc pages modified = 15).
```

```
DBCC execution completed. If DBCC printed error messages, contact a user with System Administrator (SA) role.
```

```
1> -- Confirm pages all have the correct dbid
2> SELECT 'dbcc page (Vdb5,' + convert(varchar(5), lstart) + ', 0, 0)' FROM
   master..sysusages WHERE dbid=db_id('Vdb5') AND segmap != 0
3> go

-----
dbcc page (Vdb5,0, 0, 0)
dbcc page (Vdb5,1536, 0, 0)
dbcc page (Vdb5,1792, 0, 0)
dbcc page (Vdb5,2816, 0, 0)

1> dbcc traceon(3604)
2> dbcc page (Vdb5,0, 0, 0)
3> dbcc page (Vdb5,1536, 0, 0)
4> dbcc page (Vdb5,1792, 0, 0)
5> dbcc page (Vdb5,2816, 0, 0)
6> go | grep dbid
   bmass_next=0x0 bmass_prev=0x0 bdbid=23
pageno=0 dealloc_count=102 allocnode=0 ptnid=99 allocation_page dbid=23 timestamp=000
0 00002730, segmap=0x00000003 (0x00000002
   bmass_next=0x0 bmass_prev=0x0 bdbid=23
pageno=1536 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=23 timestamp=00
00 00000001, segmap=0x00000004 (0x00000004
   bmass_next=0x0 bmass_prev=0x0 bdbid=23
pageno=1792 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=23 timestamp=00
00 00000001, segmap=0x00000003 (0x00000002
   bmass_next=0x0 bmass_prev=0x0 bdbid=23
pageno=2816 dealloc_count=0 allocnode=0 ptnid=99 allocation_page dbid=23 timestamp=00
00 00000001, segmap=0x00000003 (0x00000002
```

More Information

VDBs are unmounted and remounted during operations such as:

- Disable/Enable of the VDB.
- Stop/Start of the VDB.
- Rebooting of the server hosting the VDB's ASE instance.

DBAs are unlikely to have found the need to run **DBCC CHECKALLOC()** on physical databases in the normal workflow of resizing a database unless the ASE **MOUNT/UNMOUNT** command is used. When you **MOUNT** a database on a server that already has the mounted database's dbid in use, ASE assigns the mounted database a new dbid and this gets reflected in the **sysdatabases** and **sysusages** tables in the master database and in the **dbtable** memory structure for the mounted database. However, the dbid in the dbinfo structure and allocation pages in the mounted database are not updated by mount. Doing so can take a long time on a big database. **DBCC CHECKALLOC()** will correct the values in all the allocation pages.

In Delphix version 5.1.3.0 and higher, Delphix will create databases using a high **dbid** (using the **"CREATE DATABASE with dbid= "** clause). By using a high **dbid** (ASE supports a maximum of 32,767 dbids), it is

unlikely that the **dbid** will be reused and thus Delphix will be able to avoid the errors (14545, 14547 and 14519). It attempts to assign unique **dbid** values for staging databases in the range 10000-30000.

The second part of this fix involves keeping the staging database mounted so that the **dbid** is pinned and no other database gets created and uses the staging database's **dbid**. The staging database will remain mounted except for a brief window where we unmount it to create a manifest file.

The final part of this fix utilizes the ASE **DBCC PAGE()** command to confirm the **dbid** on select pages from the sysusages table all match the database's actual **dbid**. If there is a mismatch found, Delphix will call **DBCC CHECKALLOC()** to fix the database. Delphix calls **DBCC PAGE()** in order to avoid calling **DBCC CHECKALLOC()** as the **CHECKALLOC()** routine can add significant time to database recovery ([performance considerations for CHECKALLOC](#)).

Related Topics

- [Provisioning and Managing VDBs from SAP ASE](#)

V2P with an SAP ASE dSource or VDB

Requirements

This topic describes the procedure for exporting a virtual database (VDB) to a physical one, also known as V2P. Before performing the V2P operation, you must have created a database on the target instance into which you will load the exported data. It must be sufficiently large, and it must have been created with the **for load** SAP ASE option.

 The Delphix Engine will initiate a load command using the database specified. The V2P operation will overwrite any existing data in this database. After a V2P operation takes place, the target physical database is no longer in FOR LOAD state.

Procedure

1. Login to the **Delphix Management** application.
2. Select the **dSource** or **VDB** you want to export.
3. Select the **snapshot** of the dSource or VDB state you want to export.
4. Click **V2P**.
5. Select the **target environment**. The target environment should have been added to Delphix previously and should meet all target host requirements, see [Requirements for SAP ASE Environments and Databases](#).
6. Under **Installation**, select which **instance** that you want to export to.
7. Enter the **Name** of the database on the target instance into which you want to load the exported data.
8. Select whether or not to **Run Recovery After V2P**. When this option is set, the Delphix Engine will online the database when the export is done.
9. Click **Next**.
10. Select whether you want to have an email sent to you when the export process completes.
11. Click **Submit**.

Related Topics

- [Requirements for SAP ASE Environments and Databases](#)

Additional SAP ASE VDB Topics

- [Deleting an SAP ASE VDB](#)
- [Enabling and Disabling SAP ASE VDBs](#)
- [Migrating an SAP ASE VDB](#)
- [Refreshing an SAP ASE VDB](#)
- [Rewinding an SAP ASE VDB](#)
- [Upgrading SAP ASE VDBs](#)

Deleting an SAP ASE VDB

Procedure

 Deleting a VDB is an unrecoverable operation. Proceed only if you want to permanently destroy the unique data that was created in the VDB.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **VDB** that you want to delete.
5. From the Actions menu (...) select **Delete**.
6. If stopping or starting the VDB requires particular credentials for the target environment other than those of the default environment user:
 - a. Check **Provide Privileged Credentials**.
 - b. Enter the **username** and **password**.
 - c. Click **Validate Credentials**.
7. Click **Delete** to confirm that you want to delete the VDB.

If the VDB was currently active, the Delphix Engine will shut it down, unmount all filesystems from the target environment, and finally delete the VDB itself.

Related Topics

- [Provisioning and Managing VDBs from SAP ASE](#)
- [Enabling and Disabling SAP ASE VDBs](#)

Enabling and Disabling SAP ASE VDBs

This topic describes how to enable and disable a virtual database (VDB).

Disabling a VDB is a pre-requisite for procedures such as VDB migration or upgrade. Disabling a VDB removes all traces of it, including any configuration files, from the target environment to which it was provisioned. When you later enable the VDB again, these configuration files are restored on the target environment.

Procedure

1. Click **Manage**.
2. Select **Datasets**.
3. Click the **VDB** you want to disable.
4. From the Actions menu (...) select **Disable**.
5. Click **Disable** to acknowledge the warning.

When you are ready to enable the VDB again, from the Actions menu (...) select **Enable**, and the VDB will continue to function as it did previously.

Related Topics

- [Provisioning an SAP ASE VDB](#)
- [Deleting an SAP ASE VDB](#)

Migrating an SAP ASE VDB

This topic describes how to migrate a virtual database (VDB) from one target environment to another.

In certain situations, you may want to migrate a VDB to a new target environment. For example:

- When upgrading the host on which the VDB resides
- During a general data center migration
- To distribute the VDB load across target environments

This is easily accomplished by first disabling the database, then using the Migrate VDB feature to select a new target environment.

Prerequisites

- You must first disable the VDB before migrating it. Follow the steps outlined in [Enabling and Disabling SAP ASE VDBs](#)
- You must have already set up a new target environment that is compatible with the VDB that you want to migrate.

Procedure

1. Login to your Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **VDB** you want to migrate.
5. From the **Actions** menu (...) select **Disable**.
6. Click **Disable** to confirm.
7. From the **Actions** menu (...) select **Migrate**.
8. In the Migrate window select:
 - a. Environment: the new **target environment** for the VDB
 - b. User: select a user
 - c. Installation: the **installation** where you want to migrate
9. Click Migrate to confirm your selections.
10. From the **Actions** menu (...) select **Enable**.
11. Click **Enable** to confirm.

Within a few minutes, your VDB will re-start in the new environment, and you can continue to work with it as you would with any other VDB.

Related Topics

- [Provisioning and Managing VDBs from SAP ASE](#)
- [Enabling and Disabling SAP ASE VDBs](#)

Refreshing an SAP ASE VDB

This topic describes how to manually refresh a virtual database (VDB).

Refreshing a VDB will re-provision it from the dSource. As with the normal provisioning process, you can choose to refresh the VDB from a snapshot or a specific point in time. However, you should be aware that refreshing a VDB would delete any changes that have been made to it over time. When you refresh a VDB, you are essentially re-setting it to the state you select during the refresh process. You can refresh a VDB manually, as described in this topic, or you can set a VDB refresh policy, as described in [Managing Policies](#).

i Although the VDB no longer contains the previous contents, the previous Snapshots and Timeflow still remain in Delphix and are accessible through the Command Line Interface (CLI).

Prerequisites


To refresh a VDB, you must have the following permissions:

- **PROVISIONER** permissions on the dSource associated with the VDB
- **PROVISIONER** permissions on the group that contains the VDB
- **Owner** permissions on the VDB itself
- **Data** is a role that allows DB_ROLLBACK, DB_REFRESH, READ_ACTION, DB_SYNC, JOB_CANCEL.
- **Read** is a role that allows the user to inspect objects via the READ_ACTION permission.


A user with admin credentials can perform a VDB Refresh on any VDB in the system.

Procedure


1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **VDB** you want to refresh.
5. Click the **Refresh VDB** button.



6. Select **More Accurate** and **Next**.




7. Select the desired **refresh point** snapshot or click the **down arrow** icon to choose the Latest available range, A point in time, or An SCN to refresh from.




8. Click **Next**.

9. Click **Submit** to confirm.



10. Click the **Actions** link to watch the progress of the refresh job.

11. To see when the VDB was last refreshed/provisioned, check the **Time Point** on the Status page



Related Topics

- [Provisioning an SAP ASE VDB](#)
- [Migrating an SAP ASE VDB](#)
- [Rewinding an SAP ASE VDB](#)
- [Managing Policies](#)

Rewinding an SAP ASE VDB

This topic describes the procedure for rewinding an SAP ASE virtual database (VDB).

Rewinding a VDB rolls it back to a previous point in its TimeFlow and re-provisions the VDB. The VDB will no longer contain changes that were made after the rewind point.

i Although the VDB no longer contains changes made after the rewind point, the rolled-over snapshots and TimeFlow still remains in Delphix and are accessible through the Command Line Interface (CLI). See the topic [CLI Cookbook: Rolling Forward a VDB](#) for instructions on how to use these snapshots to refresh a VDB to one of its later states after it has been rewound.

Prerequisites

To rewind a VDB, you must have the following permission:

- **Owner** permissions on the VDB itself

You do NOT need owner permissions for the group that contains the VDB. A user with Delphix Admin credentials can perform a VDB rewind on any VDB in the system.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **VDB** you want to rewind.
5. Click the **TimeFlow** tab.
6. Select the rewind point as a snapshot or a point in time.
7. Click **Rewind**.
8. If you want to use login credentials on the target environment other than those associated with the environment user, click **Provide Privileged Credentials**.
9. Click **Yes** to confirm.

i You can use TimeFlow bookmarks as the rewind point when using the CLI. Bookmarks can be useful to:

- Mark where to rewind to – before starting a batch job on a VDB, for example.
- Provide a semantic point to revert back to, in case the chosen rewind point turns out to be incorrect.

For a CLI example using a TimeFlow bookmark, see [CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark](#).

Related Topics

- [CLI Cookbook: Rolling Forward a VDB](#)
- [CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark](#)
- [Provisioning and Managing VDBs from SAP ASE](#)

Upgrading SAP ASE VDBs

This topic describes how to upgrade an SAP ASE VDB to a higher version of SAP ASE instance.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **VDB** to be upgraded.
5. From the **Actions** menu (...) select **Disable**.
6. Click **Disable** to confirm.
7. Upgrade the ASE instance(s) hosting the VDBs.
8. Click the **Refresh** icon (under the **Manage > Environments** menu) to refresh the Delphix environment hosting the VDBs so that Delphix registers the new version of ASE.
9. Enable the VDB. The VDBs will be upgraded the first time they're enabled.
10. Repeat steps 4 to 9 for each VDB you want to upgrade.

Related Topics

- [Provisioning and Managing VDBs from SAP ASE](#)
- [Upgrading a SAP ASE dSource](#)

Backup Server Best Practices

- The backupserver log should have read permissions for the Delphix environment user
- Dump file names should be unique per source backup server.
- Dumps should be done to a local backupserver. Dumps done using the “at <backupserver>” syntax are not discovered and ingested/applied.
- Dumps of production databases that use the same backup server and same load backup path should be available on the same staging host that was pointed at during the link
- If using a remote backup server as load location, make sure “interfaces” file changes are done to ensure staging host backupserver can communicate with the remote backup server.
- A user running backupserver or process owner of the backupserver on the staging host should have read access for all the dump files that belong to dSources that use this backupserver.

Recommendations/Best Practices

1. Name the dump files as unique as possible.
2. Delphix recursively searches for backup files in a given load backup path. So assign or choose the load backup path for a given dSource to be as specific and unique as possible. For e.g, if the load backup path is /dumps for dSource db1, but the database dumps are copied to /dumps/db1, choose /dumps/db1 instead to avoid duplicate files that can happen with other dSources or databases.
3. Avoid linking sources where it's backup location is a child directory of another dSource's backup location. Ideally, each dSource using the same staging host should have its backup path set to sibling directories for ease of search.
4. Avoid copying files (of any type) with the same name as backup files that are being ingested.
5. Avoid copying other dSource or database backups to the same load backup path of the dSource of interest. You can copy, as long as dump file names are unique.
6. Only use a remote backup server as the load location if dumps cannot be available on the staging host.
7. Avoid, if possible, using production backup server as load location and the backup files are on production host when configuring a dSource with remote backupserver.
8. DE by default assumes backup files are available on the staging host when searching for them. If the backup files are available on the staging host, avoid using the remote backupserver option on the link/data management wizard.
9. If validated sync reports any fault, look at other faults that were raised, resolve all of them as well along with the current fault.
10. When copying striped backups to staging, keep all stripes in the same location.
11. If possible, create separate database dumps and/or transaction logs by more than 1 minute apart so that the dump header name is unique.
12. The Delphix Engine, by default, looks for backups that happened within the last six months from the current date/time when a source database is linked for the first time
 - a. If there is ever a need to ingest an older backup, use “Specific Backup” sync parameters and specify the location of the backup files to ingest.
 - b. Please reach out to Delphix Support for help if there is a reason for concern regarding the 6 months window.
1. The Delphix Engine periodically purges discovered backup files from its internal catalog that are older than 6 months (default). If there is a need to ingest an old backup, use the “Specific Backup” option as there is no mechanism to rediscover old backups.

Related Topics

- [SAP ASE Environments and Data Sources](#)

SAP ASE Hook Operations

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at `/bin/sh`. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh "$ARG_ENVIRONMENT_VARIABLE"
"second argument in double quotes" 'third argument in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a `bash` binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
```

```
# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

Shell Operation Tips

Using `nohup`

You can use the `nohup` command and process backgrounding from resource in order to "detach" a process from the Delphix Engine. However, if you use `nohup` and process backgrounding, you MUST redirect `stdout` and `stderr`.

Unless you explicitly tell the shell to redirect `stdout` and `stderr` in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either `stdout` or `stderr`. Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your `RunCommand` operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The `RunExpect` operation executes an `Expect` script on a Unix environment. The `Expect` utility provides a scripting language that makes it easy to automate interactions with programs which normally can only be used interactively, such as `ssh`. The Delphix Engine includes a platform-independent implementation of a subset of the full `Expect` functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Management application and CLI to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an `ssh` session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

SAP ASE Environment variables

Operations that run user-provided scripts have access to environment variables. For operations associated with specific dSources or virtual databases (VDBs), the Delphix Engine will always set environment variables so that the user-provided operations can use them to access the dSource or VDB.

dSource Environment Variables

Environment Variables	Description
ASE_ENVUSER	Environment username for the dSource
ASE_DBUSER	Database username for the dSource
ASE_DATABASE	Database name for the dSource
ASE_INSTANCE	SAP ASE Instance name for the dSource
ASE_PORT	SAP ASE Instance port for the dSource

VDB Environment Variables

Environment Variables	Description
ASE_ENVUSER	Environment username for the VDB
ASE_DBUSER	Database username for the VDB
ASE_DATABASE	Database name for the VDB
ASE_INSTANCE	SAP ASE Instance name for the VDB
ASE_PORT	SAP ASE Instance port for the VDB

Staging Server Environment Variables

Environment Variables	Description
ASE_ENVUSER	

Related Topics

- [Hook Scripts for Automation and Customization](#)

Support for Dump History File

Overview

Prior to ASE version 15.7 ESD#2, ASE did not provide any structured information about backup/restore operations performed on databases. Delphix relied on the backup server log file to find information about any backups performed using a particular backup server. The Delphix Engine parsed this log file, read each backup file to find which database the backup belonged to and restored the ones required for linked databases. As backup information is not very structured in this log file, getting backup history from this file resulted in the collection and parsing of unnecessary data. Starting with ASE 15.7 ESD#2, users can enable an option at the ASE instance level to record all backup and restore operations in a particular file, known as the Dump History file. Backup and restore operation information is stored in this file in a structured way and can be reliably used to get information about backups performed on a given database. When Dump History is enabled for a linked source, Delphix will use this feature to get information about backups for all sync operations including validated sync.

Prerequisites

To use this feature, Dump history needs to be enabled at the ASE Source instance level.

Linking A Database with Dump History Enabled

By default, Delphix uses the backup server log file to discover backup information. Users can switch to Dump History by explicitly enabling it for the dSource. Before linking a database with Dump History enabled, make sure that Dump History is enabled on the source ASE instance. To enable Dump History for the database being linked, go through the linking wizard and on the **Data Management** page, check the box marked "Use Dump History".

The screenshot shows the 'Add dSource' wizard with the 'Data Management' step selected. The configuration options are as follows:

- Initial Load:**
 - New Full Backup
 - Most Recent Existing Full Backup
 - Specific Existing Full Backups
- Backup Path:** /work
- Staging Environment @:** ASE_RH74SRC
- Repository:** ASE157_SRC
- Validated Sync Mode:** Enabled
- Use Dump History:** Enabled
- LogSync:** Enabled

Navigation buttons at the bottom: Cancel, Back, Next (highlighted), Submit.

Related Topics

- [Linking Data Sources and Syncing Data with SAP ASE](#)

IBM Db2 Environments and Data Sources

In this section, we will describe all the different Delphix and Db2 components that are part of the Delphix Virtualization solution for IBM Db2 database.

This section covers the following topics:

- [IBM Db2 Overview](#)
- [Delphix Architecture with Db2](#)
- [Db2 Plugin Installation and Upgrade](#)
- [How to use dlp_x_db_exec privilege elevation script](#)
- [Quick Start Guide for Db2](#)
- [Db2 Support and Requirements](#)
- [Managing Db2 Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with Db2](#)
- [Provisioning and Managing VDBs from Db2 dSources](#)
- [Db2 Hook Operations](#)
- [Enhanced Logging for Db2 Plugin](#)

IBM Db2 Overview

Introduction to Db2

Db2 for Linux, UNIX, and Windows (LUW) is a database server product developed by IBM. Db2 LUW is the "Common Server" product member of the Db2 family, designed to run on most popular operating systems.

The version numbers in Db2 are non-sequential with v11.1 and 11.5 being the two most recent releases. Specifics of Db2 versions and platforms supported on Delphix are located in the [Db2 Compatibility Matrix](#).

IBM Db2 Authentication

Db2 User and group authentication is managed in a facility external to Db2 LUW, such as the operating system, a domain controller, or a Kerberos security system. This is different from other database management systems (DBMSs), such as Oracle and SQL Server, where user accounts may be defined and authenticated in the database itself, as well as in an external facility such as the operating system.

Any time a user ID and password is explicitly provided to Db2 LUW as part of an instance attachment or database connection request, Db2 attempts to authenticate that user ID and password using this external security facility. If no user ID or password is provided with the request, Db2 implicitly uses the user ID and password that were used to login to the workstation where the request originated. More information on Db2 authentication and authorization is available via [IBM documentation](#).

Delphix Db2 Authentication

Delphix for Db2 requires that the staging and target hosts must already have the necessary users and authentication systems created/installed on them. Delphix will neither create users nor change database passwords as part of the provisioning process.

Db2 Database Level Support

Delphix supports a number of Db2 database-level features. An overview of Db2 database level support is as follows:

- Support for multiple databases linking in a Single Instance, which allows Delphix Engine users to utilize an available instance on the target more efficiently.
- Support for using a customer-supplied directory for the Delphix Plugin, Db2 Mount Points, and Db2 Delphix files and logs.
- Support for the in-memory BLU feature of Db2.
- Support for Kerberos environments.
- Support for VDB provisioning of the same OS version level or one version higher than the source and staging instances.
- Intelligent handling of HADR logs.

High Availability Disaster Recovery (HADR)

The HADR feature of IBM Db2 provides a high availability solution for both partial and complete site failures. It protects against data loss by replicating data changes from a source database, called the primary, to one or more target databases, called the standby.

Delphix HADR Support

HADR replication takes place at a database level, not at the instance level. Therefore, a standby instance can have multiple databases from multiple different primary servers/instances on it. If the instance ID on the Delphix standby is not the same as the instance ID on the primary, the Delphix standby instance ID must have database permissions **secadm** and **dbadm** granted to it on the primary database. These permissions and all HADR settings must be implemented on the primary database before you take the backup on the primary database.

Log Transmitting

All changes that take place at the primary database server are written into log files. The individual log records within the log files are then transmitted to the secondary database server, where the recorded changes are replayed to the local copy of the database. This procedure ensures that the primary and the secondary database servers are in a synchronized state. Using two dedicated TCP/IP communication ports and a heartbeat, the primary and the standby databases track where they are processing currently, the current state of replication, and whether the standby database is up-to-date with the status of the primary database. When a log record is "closed" (still in memory, but has yet to be written to disk on the primary), it is immediately transmitted to the HADR standby database(s). Transmission of the logs to the standbys may also be time-delayed.

Multiple Standby

Beginning in Db2 v10.1, the HADR feature supports multiple standby databases. This enables an advanced topology where you can deploy HADR in multiple standby mode with up to three standby databases for a single primary. One of the databases is designated as the principal HADR standby database, with the others termed as auxiliary HADR standby databases. As with the standard HADR deployment, both types of HADR standbys are synchronized with the HADR primary database through a direct TCP/IP connection. Furthermore, both types support the reads on standby feature and can be configured for time-delayed log replay. It is possible to issue a forced or unforced takeover on any standby, including the delphix auxiliary standby. However, you should never use the Delphix auxiliary standby as a primary, because this will impact Delphix performance.

Delphix HADR Synchronization

The Delphix for Db2 uses the HADR capability of Db2 to synchronize data from a production Db2 database into a Delphix-controlled Db2 "standby" server. By using this mature and existing Db2 capability, the Delphix Engine is able to ingest data and keep the standby server in sync with only a minimal impact on production. The HADR connection is configured to Super-Asynchronous (SUPERASYNC) mode where log writes are considered successfully transmitted when the log records are sent from the primary database. Because the primary database does not wait for acknowledgments from the standby database, there is no delay on the primary and transactions are considered committed regardless of the state of the replication of that transaction. For further information on Delphix synchronization, see [Linking a dSource from a Db2 Database: An Overview](#).

Database Partitioning Feature

Database Partitioning Feature (DPF) lets you partition your database across multiple servers. Since you can add new machines and spread your database across them, this allows users to scale their database. This means more CPUs, more memory, and more disks from each of the additional machines are available for your database. DPF can be used to manage large databases for a variety of use cases including data warehousing, data mining, online analytical processing (OLAP), or online transaction processing (OLTP) workloads.

DPF enables the user to divide a database into database partitions, a database partition is a part of a database that consists of its own data, indexes, configuration files, and transaction logs. Each database partition can be configured on the different physical server having its own set of computing resources, including CPU and storage. When a query is processed, the request is divided so each database partition processes the rows that it is responsible for. DPF can maintain consistent query performance as the table grows by providing the capability to

add more processing power in the form of additional database partitions. This capability is often referred to as providing linear scalability using Db2s shared-nothing architecture.


DPF is an approach to sizing and configuring an entire database system. Please follow the recommended practices for optimal performance, reliability, and capacity growth. Please refer to IBM documentation of DPF for more details in [IBM knowledge center](#).

Related Topics

- [IBM Db2 Environments and Data Sources](#)

Delphix Architecture with Db2


This topic describes the high-level process for adding Db2 environments, linking Db2 instances to the Delphix Engine, and provisioning virtual databases.



Delphix uses a Standby server called “Staging” to support the data ingestion process. There are multiple benefits of having a Staging server between the Db2 Source Instances and the Delphix Dynamic Data Platform:

- Allows the database instances to stay in sync with the source database, so end users can create or update their virtual databases as needed. The Delphix Engine snapshots the Staging server and the snapshots can be provisioned out to one or more target servers.
- Supports multiple ways of ingesting data, via HADR, Backups, Archive Logs, and Staging Push.
- In the case of the Staging Push feature, the user will control the restore and roll forward of the staging database.
- Prevents the end-users and their related processes to be connected to the primary database instances, making the architecture more secure and reducing primary databases' potential point of failure and performance degradation.
- In the case of Database Partitioning Feature (DPF), the staging and target environments will be single-node environments and all the portions on the source (logical and physical) will be converted to logical partitions on a single host.

Database Partitioning Feature Architecture



⚠ The snapshot and provision process occurs on the instance level; all databases that exist on the standby server will be provisioned out to the target machines. Similarly, actions such as bookmark, rewind, and refresh will simultaneously apply to all the databases in the instance.

Related Topics

- [IBM Db2 Environments and Data Sources](#)

Db2 Plugin Installation and Upgrade

The Db2 plugin is provided as a zip file on the Delphix download site. Follow the steps below to download the plugin.

1. In the web browser, go to the Delphix [download](#) site.
2. Login to the download site using email and password credentials.
3. Navigate to the required version number of Delphix Engine.
4. Go to the Plugins > Db2 > Plugin_<version_number>.zip folder and download the zip file.


For initial installation and upgrading to a higher version refer to the section [Delphix Engine Plugin Management](#) for further details. For upgrading the Lua-based plugin to the vSDK plugin, refer to the following section.

Upgrading the Lua-based Plugin to the vSDK-based Plugin

The steps to upgrade from the Lua-based plugin to the new vSDK-based plugin are described below.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. You need to disable all the datasets in the Lua environment before upgrading.
Click the dSource you want to disable.
5. From the **Actions** menu (...) select **Disable**.
6. Repeat the above step to disable all the VDB datasets.
7. Select **Manage > Plugins**.
8. To upload a vSDK Plugin, click the  icon.
This opens the Upload or Upgrade a Plugin dialog. Upload the <file_name>.json vSDK plugin.



9. To start using the plugin, refresh target and source environments to re-discover repositories.
You will notice the changes in the toolkit directory structure as compared to the Lua plugin environment.

 This plugin upgrade operation will read the `advanceConfFlag.conf` and `plugin-logging.properties` files present in the Lua plugin environment and combines them to create `db2_plugin.conf`. This operation also creates the `plugin_3.x_archives` directory and moves all the Lua-specific files such as the logs directory, the code directory, and the configuration files (`advanceConfFlag.conf` and `plugin-logging.properties`) into this directory. The plugin archive directory is for reference purposes only. You can delete the same if not required.

10. Click **Manage**.
11. Select **Datasets**.
12. From the **Actions** menu (...) select **Enable**.
This will enable the existing datasets under the new vSDK environment.
13. The upgrade is now complete. You can now check the logs for upgrade-related information.

Related Links

- [Db2 Compatibility Matrix](#)
- [Delphix Engine Plugin Management](#)

How to use dlpx_db_exec privilege elevation script

The following section describes how the Db2 Plugin works with a Kerberos environment. Customers use principal (A user in Kerberos is called a principal) to connect with DE. All operations where Plugin needs to execute commands using database users, in that case, privilege elevation script will be used.

The following example shows how the Db2 Plugin works in the Kerberos enabled environment for a customer:

1. All SSH connections to remote hosts use a single environment user that is principal in the case of a Kerberos environment.
2. All Db2 commands which need to be executed using a particular OS user (in the case of DB2 it'll be an instance user) will pass through a privilege elevation script. Commands requiring privilege elevation will be executed under a script dlpx_db_exec, with the first parameter being the user to execute as, and the remaining parameters being the command to execute. This script may be customized by the customer, but it must always return the results of the executed command and exit with the return code from the executed command.
3. When a command is invoked with dlpx_db_exec then we'll pass the instance IDs which was discovered during the discovery phase.

Implementation

1. The customer will specify the same primary environment user for the execution of all Delphix object requests.
2. All SSH connections will then be made with this user.
3. A new utility dlpx_db_exec has been added to the Plugin.
 - a. The first param to this command will be the db2 instance user to execute the remaining args.
 - b. This script will be customizable by the customer to use their internal utility or command to essentially sudo to the requested instance user.
1. The Db2 Plugin is updated to pass the instance ID string to the dlpx_db_exec script such that the customer can update this to use sudo (or some other custom elevation utility).
2. The Db2 Plugin scripts are updated to prefix all commands required to be executed as the instance ID with dlpx_db_exec.
3. All Db2 changes are based on top of the Db2 DB level changes.

Sample content of dlpx_db_exec script

```
#
# This script allows customization of command execution with an alternate user
# account.
# Arg $1 contains "-u<optional user account>" for the desired user under
# which database commands will be executed.
# By default this argument is ignored and the script is executed as the default
# account.
```

```
#

if [[ $1 != -u* ]]; then

    echo "Incorrect command line parameters, -u<optional user account> is
    required as the first parameter"

    exit 1

fi

user_id=`echo $1 | sed -e "s/^-u//"`

echo "$user_id" >> /tmp/test.log

shift 1

echo "$user_id and $DB2_DB_NAME" >> /tmp/test.log

if [[ $user_id != "delphix_os" ]]; then

command=$(printf "%s " "$@" )

sudo su - $user_id -c "cd /home/delphix_os;export DB2DBDFT=$DB2DBDFT;$command"

else

$@

fi
```

Related Topics

- [Provisioning and Managing VDBs from Db2 dSources](#)

Quick Start Guide for Db2

This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with DB2 database objects in the Delphix Engine. It does not cover any advanced configuration options or best practices, which can have a significant impact on performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix functionality. It assumes you will use the VMware Hypervisor.

Overview

In this guide, we will walk through deploying a Delphix Engine, starting with configuring Source and Target database environments. We will then create a dSource, and provision a VDB.

For purposes of the QuickStart, you can ignore any references to Replication or Masking.

Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to the Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online. This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
                    example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
                    the primary interface. Setting this value
                    to "true" will cause all other properties (address,
                    hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
                    "1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
                    for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
                    -- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false

```

```
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24
```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

8. Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`)

9. Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

10. Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the browser at:
```

```
http://<address>
```

Type "exit" to disconnect, or any other commands to **continue** using the CLI.

11. Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
12. Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at <http://<Delphix Engine>/login/index.html#serverSetup>.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (sysadmin) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (admin) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

 The default domain user created on Delphix Engines is now **admin** instead of delphix_admin. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

System Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications.

Choose your option to set up system time in this section.

For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click Next to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.
4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

 Although your Delphix Engine will work without registration, we strongly recommend that you register each Delphix Engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

Db2 Source and Target Environment Requirements

Source Db2 hosts are servers that have Db2 binaries installed and have Db2 instances created on them. The hosts that contain the data that we wish to ingest are referred to as the source environment. Hosts with empty instances (no dbs in instance) are used as either staging or target hosts. This topic describes the requirements for creating connections between the Delphix Engine and Db2 hosts and instances.

Requirements for Db2 Source Hosts and Instances

Each Db2 Source host must have IBM Db2 installed and at least one instance created on the machine. Depending on the way you intend to ingest data into the Staging server, you will also need to ensure that the following requirements are met:

- For HADR Ingestion: HADR settings for each database to be used with the standby server should be preset before the linking process begins as described in [Linking a Db2 dSource](#)
- For Backup and Log Ingestion, no additional requirements are needed at this point.

Requirements for Db2 Staging and Target Hosts and Instance

- The staging environment that the Delphix Engine uses must have access to an existing full backup of the source database on disk to create the first full copy. Delphix recommends using compressed backups as that will reduce storage needs and speed up ingest. Once this first full backup is available, subsequent refreshes can be accomplished via HADR or Archive Logs.
- The staging and target Db2 instances that you wish to use must already exist on the host. We can use the same instance for the dSource and VDB creation and can also have multiple VDBs on the same instance.
- The available instances on each host can be verified by going to the databases tab for the environment in question.
- Instance level configuration values such as the bufferpool value will need to be managed by the customer independent of Delphix.
- The instances used for staging and target environments must be compatible with the source Db2 instance.

Requirements for Db2 DPF Staging and Target Hosts and Instances

 It is highly recommended that the Database Partitioning Feature (DPF) for Db2 staging and target should be configured on separate hosts due to reason that the DPF environment consumes a lot of resources.

- The staging and target hosts should be configured with the same number of logical partition nodes as the source.
- The logical portion configuration should be added in \$HOME/sql/lib/db2nodes.cfg file where \$HOME is the home directory of the instance user.
- The hostname to be used should be a fully qualified domain name.
- db2 nodes.cfg file configured for 4 logical partitions should look like:

```
0 <FQDN hostname> 0
1 <FQDN hostname> 1
2 <FQDN hostname> 2
```

```
3 <FQDN hostname> 3
```

- Each logical partition should have a corresponding port entry in `/etc/services` of hosts. In the example below, `db2inst1` is the instance the user and port number from 60000 to 60003 are reserved for inter-partition communication.

```
DB2_db2inst1 60000/tcp
DB2_db2inst1_1 60001/tcp
DB2_db2inst1_2 60002/tcp
DB2_db2inst1_END 60003/tcp
```

- Registry variables `DB2RSHCMD` and `DB2COMM` should be updated with values as shown below:

```
SSHCMD=$(which ssh)
db2set DB2RSHCMD="$SSHCMD"
db2set DB2COMM=TCPIP
```

- After the environment is configured on staging and target, execute **db2_all echo hello** on both staging and target servers to add host key in the `known_hosts` file. This will also help to check for any ssh issues on these servers.
- The backup files (or log directory if using **Customer Supplied Archive Logs**) should be consistent with the partitions configured on staging.

Additional Environment Requirements

- There must be an operating system user (`delphix_os`) with these privileges:
 - Ability to login to the staging and target environment via SSH
 - Ability to run `mount`, `umount`, `mkdir`, and `rmdir` as a super-user. If the staging and target host is an AIX system, permission to run the `nfso` command as a super-user. See [Sudo Privilege Requirements for Db2 Environments](#) for further explanation of the commands and [Sudo File Configuration Examples for Db2 Environments](#) for examples of the `/etc/sudoers` file on different [Outdent](#) operating systems.
- There must be a directory on the staging and target environment where you can install the Delphix Engine Plugin. For example, `/var/opt/delphix/toolkit`.
 - The `delphix_os` user must own the toolkit path directory (`__var/opt/delphix/toolkit`).
 - If the `delphix_os` user and instance users (responsible for the Delphix Engine operations such as linking and provisioning) are sharing any common group, then the toolkit path directory (`__var/opt/delphix/toolkit`) must have permissions `rwxrwx-- (0770)`, but you can also use more permissive settings. This allows instance users who are part of the same group as the `delphix_os` user's group to be able to create directories inside the toolkit path directory.
 - If the `delphix_os` user and instance users (responsible for the Delphix Engine operations such as linking and provisioning) are not sharing any common group, then the toolkit path directory must have `-rwxrwxrwx (0777)` permissions.
 - The directory should have 1.5GB of available storage: 400MB for the Plugin and 400MB for the set of logs generated by each Db2 instance that runs on the host
 - In Db2 Plugin, plugin directory space will be used as the base location for the mount point
 - When multiple engines access the same host environment, create a separate Db2 plugin directory for each engine that will discover this environment.

 For Db2 plugin version 4.0.0 and onwards, follow the steps below to change the toolkit directory.

- i. Disable all the datasets related to that environment
- ii. The instance of the dataset must have the DIR_CACHE configuration parameter set to No, as shown below.
db2 update dbm cfg using DIR_CACHE NO
- iii. Upload the new vSDK-based plugin. The plugin has upgrade logic implemented in it and will upgrade the existing Lua-based plugin.
- iv. Refresh the environment to allow the plugin to create plugin paths and files under the scratch path.
- v. Enable all the datasets and complete the migration process.

- The Delphix Engine must be able to initiate an SSH connection to the staging and target environment
- NFS client services must be running on the staging and target environment
- Staging and Target host instances must have the same Db2 instance code page as the production instance. The instance owner must have a compatible value defined in LC_ALL or LC_CTYPE for the Db2 code page. Preferably derived from the value of LC_ALL, LC_CTYPE, or LANG of the instance owner on the production host.
- The primary environment user should share its primary group with the instance user. For example, if the delphix_os is the primary environment user which is used for environment addition and its primary group is also delphix_os then instance users (responsible for the Delphix Engine operations such as linking and provisioning) should share group delphix_os as their secondary group.

Custom Configurations in Db2 Plugin

DB2 Plugin v3.x

The advanced config file is named **advanceConfFlag.conf** and is created under the **<Toolkit directory>** during discovery or environment refresh operation if the file does not pre-exist already. This file will have the following configurable parameters:

- **Common Group (*notCommonGroupFlag*):**
 - This parameter sets the permission of "mnts", "logs" and "code" directories under the **<Toolkit dir>/DB2** location.
 - Set ***notCommonGroupFlag*** to *false*, if the primary group of the primary environment user (user which is used to do discovery) is shared with the instance user.
 - Set ***notCommonGroupFlag*** to *true*, if the primary environment user and the instance user do not share any user group.
 - Once the necessary changes are made, refresh the environment in order to apply the changes.

By default, the ***notCommonGroupFlag*** parameter is commented out in the file. This means that the plugin implicitly assumes its value as false. The valid values for this parameter are true or false.

- **Allow Source Database on Same Instance (*allowSourceDBonSameInst*)**

By default, the DB2 Plugin restricts the provisioning of a VDB on a DB2 instance which contains a database with the same name as the VDB's source database

For example:

- When provisioning a VDB from a VDB then the latter is treated as the source database.
- When provisioning a VDB from a dSource then the staging database representing the dSource is treated as the source database.

Set ***allowSourceDBonSameInst*** to true, if the user wants to provision a VDB on an instance which contains a database with the same name as the VDB's source database. The valid values for this parameter are true or false.

- **Restore Pipeline limit (*restorePipelineLimit*):**

Db2 DPF Plugin performs restoration of all non-catalog partitions in parallel.

The parameter ***restorePipelineLimit*** allows the user to configure the number of partitions that could be restored in parallel by the plugin. When the new **advanceConfFlag.conf** file is created during environment refresh or discovery operation, the default value of ***restorePipelineLimit*** is set to 10. The user can modify the value of this parameter to a desired value. The valid value for this field is any positive integer.

A default config file created during discovery or environment refresh operation:

```
# If the user is not sharing a common group between primary environment user and
instance user.

# Then the user needs to set parameter notCommonGroupFlag as true.

# notCommonGroupFlag=true

# During provision operation check if instance contains a database name which is
identical to the source DB name used for provisioning operation.

# By default plugin will never allow creating a VDB on the instance where source
database (or other database which is identical in name with source DB) already
exists.

# If the user still wants to create a VDB on the instance which contains a database
name which is identical to the source DB name used for provisioning operation,

# then the user needs to set parameter allowSourceDBonSameInst as true.

# allowSourceDBonSameInst=true

# Limit for parallel restores

restorePipelineLimit=10
```

DB2 Plugin v4.x and onwards

DB2 plugin config file is named **db2_plugin.conf** and is created under the **<Toolkit directory>/<Delphix_COMMON>/plugin/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/** directory during discovery or environment refresh operation if the file does not already exist. This file has two sections represented by the headers - `plugin_custom_parameters` and `plugin_logging_parameters`.

 It is important for the users to maintain two section headers in the `db2_plugin.conf` file. If any of the sections get deleted, the plugin will pick the default value for the parameters under the deleted section.

Plugin Custom Parameters

Common Group (*usersHaveCommonGroup*)

- This parameter sets the permission for the "mnts" directory located under the **<Toolkit dir>/DB2** directory and for the "logs" directory located under the **<Toolkit directory>/<Delphix_COMMON>/plugin/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/** directory.
- Set the **usersHaveCommonGroup** parameter as **True**, if the primary group of the primary environment user (user which is used to do discovery) is shared with the instance user.
- Set the **usersHaveCommonGroup** parameter as **False**, if the primary environment user and the instance user do not share any user group.
- Once the necessary changes are made, refresh the environment in order to apply the changes.

By default, the **usersHaveCommonGroup** parameter is set to **True**. The valid values for this parameter are true or false.

Allow Source Database on Same Instance (*allowSourceDbOnSameInstance*)

By default, the DB2 Plugin restricts the provisioning of a VDB on a DB2 instance which contains a database with the same name as the VDB's source database

For example:

- When provisioning a VDB from a VDB then the latter is treated as the source database.
- When provisioning a VDB from a dSource then the staging database representing the dSource is treated as the source database.

If the user wants to provision a VDB on an instance which contains a database with the same name as the VDB's source database, then the parameter **allowSourceDbOnSameInstance** must be set to true. The valid values for this parameter are true or false.

Restore Pipeline limit (*restorePipelineLimit*)

Db2 DPF Plugin performs restoration of all non-catalog partitions in parallel.

The parameter **restorePipelineLimit** allows the user to configure the number of partitions that could be restored in parallel by the plugin. When the new **advanceConfFlag.conf** file is created during environment refresh or discovery operation, the default value of **restorePipelineLimit** is set to 10. The user can modify the value of this parameter to the desired value. The valid value for this field is any positive integer.

Plugin Logging Parameters

This feature provides three options to configure and control enhanced logging by the plugin:

1. To set debug log levels.
2. To set active log file size.
3. To set log retention levels.

Delphix users can modify the above parameters defined under the `plugin_logging_parameters` section of this file.

Setting Log Levels (level)

Delphix users can set the logging level of the plugin-generated logs by setting the parameter **level** under the `plugin_logging_parameters` header. Higher logging levels will help to expedite debugging issues. There are two levels of logging which are: Info and Debug. The configured logging level will apply to all the objects (both dSources and VDBs) present on a staging/target host. For example; if we have 4 dSources on a host associated with a Delphix Engine then the log level will apply to all the dSources.

Log level description:

- level=INFO: This level will print only informational logs. This is the default log level.
- level=DEBUG: This level will print informational logs and debug statements.

Setting Maximum Size of Active Log File (*logFileSize*)

Delphix users can set the maximum size of the active log file by defining the value of the `logFileSize` parameter under the `plugin_logging_parameters` header of the `db2_plugin.conf` file. Once this maximum size limit is reached, the plugin will rotate the log as per the retention property described below. This parameter only takes in a positive integer value. The minimum value of this parameter is 1 MB and the maximum value is 10 MB. The default value is 1 MB.

Setting Log Retention Levels (*retention*)

Delphix users can set a retention level for plugin-generated logs for each dataset (`<Db name>.diag.log`) using the parameter ***retention*** under the `plugin_logging_parameters` header. As per this parameter, the log files are moved (archived), renamed or deleted once they reach the value set in the `logFileSize` parameter. New incoming log data is directed into a new fresh file (at the same location).

By default, the value for this parameter will be set to a minimum value of 2. The user can change this value and set its value within the range 2 and 50. For example; if ***retention*** is set to 4, the plugin will have the following log files: `<DB Name>.diag.log`, `<DB Name>.diag.log.1`, `<DB Name>.diag.log.2`, `<DB Name>.diag.log.3`, `<DB Name>.diag.log.4`.

- File `<DB Name>.diag.log` is the active log file
- File `<DB Name>.diag.log.1` is the most recent archive log file
- File `<DB Name>.diag.log.4` is the oldest one

 The DB2 plugin will ignore values for a parameter if the values do not fit the description above and will continue the operation with the default set of values for that particular parameter.

db2_plugin.conf

```
#
# The parameters must not be left padded with spaces. The plugin will otherwise not
# accept the specified parameter
# values and will use the default values.
#

[plugin_custom_parameters]
#
# If the user has not shared a common group between primary environment user and
# instance user.
# Then the user needs to set parameter usersHaveCommonGroup as False.
#

usersHaveCommonGroup=True

#
# During provision operations the plugin checks if the target instance contains a
# database name which is identical to
# the source DB name from which the VDB is to be created. By default the plugin will
# never allow creating a VDB on the
# instance where the source database already exists, or on any instance that has a
# database with the same name. If the
# users still want to create a VDB on an instance which contains a database with the
# same name as the source DB then
# they need to set parameter allowSourceDbOnSameInstance as True.
```

```
#  
  
allowSourceDbOnSameInstance=False  
  
#  
# If the users want to control the number of parallel restores then they can tune the  
# parameter restorePipelineLimit.  
# By default the value of this parameter would be 10.  
#  
  
restorePipelineLimit=10  
  
[plugin_logging_parameters]  
#  
# This flag will set the debug level of plugin logs on the remote server (both  
# staging and target). There are two levels:  
# Info  
# Debug  
# The above are the only valid values that can be assigned. If any other value is  
# assigned, the plugin will set the  
# default level as Info. the string is case insensitive so info, Info and INFO are  
# acceptable.  
#  
  
level=INFO  
  
#  
# This parameter will set the maximum size of the active log file in MB. Once this  
# limit is reached, the plugin will  
# rotate the log as per the retention property defined below. The minimum value of  
# this parameter is 1 MB and maximum  
# value is 10 MB. This parameter only takes in a positive integer value. The default  
# value is 1 MB  
# logFileSize=<positive integer>  
#  
  
logFileSize=1  
  
#  
# Whenever the size of <DB Name>.diag.log file exceeds the value provided by  
# logFileSize parameter then the plugin will  
# rename the active log file to <DB Name>.diag.log.<number> and a new log file with  
# name <DB Name>.diag.log will be  
# generated. For example, if LogRetention is set to 4, the plugin will have the  
# following log files: <DB Name>.diag.log,  
# <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB Name>.diag.log.3, <DB  
# Name>.diag.log.4.  
# File <DB Name>.diag.log.4 will be the oldest one.  
# File <DB Name>.diag.log.1 will be the most recent archive log file.  
# File <DB Name>.diag.log will be the active log file.  
# The minimum and default value for retention flag is 2.  
# retention=<positive integer>
```

```
#
retention=2
```

Instance User Requirements

- The instance owner of each instance you wish to use within staging or a target host must be added as an environment user within the Delphix engine. See [Managing Db2 Users and Instance Owners](#).
- For HADR synced dSources the staging instance owner must be able to "read" the ingested database contents as Delphix will check the validity of the database by querying tables before each dSource snapshot.

Database Container Requirements

- All Db2 database container types are fully supported with the exception of Db2 raw containers.
- It is recommended not to use special characters for source database container paths as this may create problems while parsing such container paths and may also result in failures during dSource and VDB creation.

 If a container is added or deleted, the dSource will have to be resynced.

Sudo Privilege Requirements for Db2 Environments

This topic describes the rationale behind specific `sudo` privilege requirements for virtualizing Db2 Databases.

Privilege	Sources	Targets and Staging	Rationale
<code>mkdir/rmdir</code>	Not Required	Required	Delphix dynamically makes and removes directories under the provisioning directory during VDB operations. This privilege is optional, provided the provisioning directory permissions allow the delphix os user to make and remove directories.
<code>mount/umount</code>	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for superuser.

 It is required to specify the NOPASSWD qualifier within the "sudo" configuration file, as shown here: [Sudo File Configuration Examples for Db2 Environments](#). This ensures that the "sudo" command does not demand the entry of a password, even for the "display permissions" (i.e. "sudo -l") command. Delphix issues "sudo -l" in some scripts to detect if the operating system user has the correct sudo privileges. If it is unable to execute this command, some actions may fail and Delphix will raise an alert suggesting it does not have the correct sudo permissions. Restricting the execution of "sudo -l" by setting "listpw=always" in the "/etc/sudoers" file when the Delphix operating system user is configured to use

public key authentication will cause the Delphix operating system user to be prompted for a password which will fail certain Delphix actions. Use a less restrictive setting for listpw than "always" when the Delphix operating system user is using public-key authentication.

Adding a Db2 Environment

Prerequisites

- Make sure that the staging environment in question meets the requirements described in [Requirements for Db2 Hosts and Databases](#)

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions (...) menu and** select **Add Environment**.
5. In the Add Environment dialog, select **Unix/Linux** in the menu.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter a **Name** for the environment.
9. Enter the **Host IP address** or **hostname**.
10. Enter the **SSH port**.
The default value is 22.
11. Enter an **OS Username** for the environment.
For more information about the environment user requirements, see [Requirements for Db2 Hosts and Databases](#).
12. Select Login Type.
 - Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

 Currently, the Enterprise Password Vault option is not supported.

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/ .ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For Password, enter the password associated with the user in step 11.
14. Using Public-Key Authentication:
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write privileges to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.
15. You can also add public key authentication to an environment user's profile by using the command-line interface, as explained in the topic [CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users](#).
16. For Password Login, click Verify Credentials to test the username and password.
17. Enter a **Toolkit Path** (make sure the toolkit path does not have spaces).
For more information about the toolkit directory requirements, see [Requirements for Db2 Hosts and Databases](#).
18. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
19. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
20. Click **Submit**.

As the new environment is added, you will see two jobs running in the Delphix platform Job History, one to Create and Discover an environment, and another to Create an environment. When the jobs are complete, you will see the new environment added to the list in the Environments tab. If you do not see it, click the Refresh icon in your browser.

Managing DB Instances

When you add an environment with the Delphix Management Application, all existing DB2 instances on the host are automatically discovered by Delphix. A list of all instances and databases are available to Delphix based on the environment discovery process.

View Instances

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the Environments panel, click on the **name of the environment** you want to refresh.
5. Select the **Databases** tab to see a list of all DB2 instances found in the environment.

Linking a Data Source (dSource)

Prerequisites

- Be sure that the source and staging instances meet the host requirements and the databases meet the container requirements described in [Requirements for Db2 Hosts and Databases](#).

Source Database Preparation

- Instance Owner Permissions

Delphix uses the Db2 instance owner account on the dSource for multiple operations, including verifying the data inside the databases. In order to ingest a database on the staging server with different instances, we need permissions on the source database in order to perform a restore on the staging server. As an example, if we have an instance named "auto1051" at the source with a database name "delphix" and we want to create a dSource on the "auto1052" instance on the staging server, then you must explicitly grant DBADM

and SECADM to the dSource instance "auto1052" on the source instance using the following steps:

- a. Connect to the source database as the source instance owner.
 - i. *connect to <DB_NAME> user <SOURCE_INSTANCE_OWNER>*
- b. Issue database grant command
 - i. *grant DBADM, SECADM on the database to user <DSOURCE_INSTANCE_OWNER>*
 - ii. *grant DBADM, SECADM on the database to user <VDB_INSTANCE_OWNER>*
- c. Repeat step 2 for every database to be included in the dSource, on the corresponding source database.

Determine if your dSource will be a non-HADR instance, a HADR single standby instance, or a HADR multiple standby instance. Non-HADR dSources can only be updated via a full dSource resync from a newer backup file

Non-HADR Database

- a. See the "InstanceOwner Permissions" section above.
- b. Uncheck the default snapsync policy during the dsource linking operation
- c. Ensure that the source database has the necessary user permissions for the provisioned VDBs as described in [Database Permissions for Provisioned Db2 VDBs](#).

HADR Single Standby Database

- a. All items in Non-HADR Database section above.
- b. The following database configuration settings must be set:
 - i. *update db cfg for <DB_NAME> using HADR_LOCAL_HOST <PRIMARY_IP> HADR_LOCAL_SVC <PRIMARY_PORT> immediate*
 - ii. *update db cfg for <DB_NAME> using HADR_REMOTE_HOST <STANDBY_IP> HADR_REMOTE_SVC <STANDBY_PORT> immediate*
 - iii. *update db cfg for <DB_NAME> using HADR_REMOTE_INST <STANDBY_INSTANCE_NAME> immediate*
 - iv. *update db cfg for <DB_NAME> using HADR_SYNCMODE SUPERASYNC immediate*
- c. If database configuration parameter LOGINDEXBUILD is set to OFF, do the following:
 1. *update db cfg for <DB_NAME> using LOGINDEXBUILD ON*
 2. Force off all connections to the database and reactivate the database
- d. If database configuration parameter LOGARCHMETH1 is set to OFF, do the following:
 - i. *update db cfg for <DB_NAME> using LOGARCHMETH1 XXXX (must be a valid log archiving method)*
 - ii. Take an offline backup
- e. If LOGARCHMETH1 points to a third-party backup server (i.e. TSM or Netbackup) define LOGARCHMETH2 to disk
 - i. *update db cfg for <DB_NAME> using LOGARCHMETH2 DISK:<full path to archive log directory>*
 1. Log files in the directory must be available from the time of the backup until the restore has successfully completed on the dSource.
- f. *db2 start hadr on db <DB_NAME> as primary by force*
- g. Take a full online backup as defined in the "Backup Source Database" section below.
- h. Record the following information, as it must be entered on the Delphix Engine while creating the dSource.
 - i. HADR Primary hostname
 - ii. HADR Primary SVC
 - iii. HADR Standby SVC (auxiliary standby port)

HADR Multiple Standby Databases

This **assumes** a single standby database HADR setup already exists. The existing standby will be referred to as the main standby. The new delphix standby will be referred to as the auxiliary standby.

- a. The following database configuration settings must be set on the primary database:
 - i. *update db cfg for <DB_NAME> using HADR_SYNCMODE <SYNC MODE> immediate* – set whichever sync mode you wish to use on your main standby.
 - ii. *update db cfg for <DB_NAME> using HADR_TARGET_LIST "<MAIN_STANDBY_IP:MAIN_STANDBY_PORT|AUXILIARY_STANDBY_IP:AUXILIARY_STANDBY_PORT>" immediate*
 1. You may have up to two auxiliary standbys defined separated by a '|'; one of which must be the delphix dSource.
- b. *stop hadr on db <DB_NAME>*
- c. *start hadr on db <DB_NAME> as primary by force*
- d. Take a full online backup as defined in the "Backup Source Database" section below. While this backup is running, you may continue with step 5.
- e. The following database configuration settings must be set on the existing main standby database:
 - i. *update db cfg for <DB_NAME> using HADR_SYNCMODE <same mode as defined in 1.a above.>* – It must be the same value used for the primary database.
 - ii. *update db cfg for <DB_NAME> using HADR_TARGET_LIST "<PRIMARY_IP:PRIMARY_PORT|MAIN_STANDBY_IP:MAIN_STANDBY_PORT>"*
- f. *stop hadr on db <DB_NAME>*
- g. *start hadr on db <DB_NAME> as standby*
- h. Record the following information, as it must be entered on the Delphix Engine while creating the dSource (the auxiliary standby database):
 - i. HADR Primary hostname
 - ii. HADR Primary SVC
 - iii. HADR Standby SVC (auxiliary standby port)
 - iv. HADR_TARGET_LIST <PRIMARY_IP:PRIMARY_PORT|MAIN_STANDBY_IP:MAIN_STANDBY_PORT>

Delphix Recovery with Customer Supplied Archive Logs

The Delphix Recovery with Customer Supplied Archive Logs (formerly known as Delphix Backup and Logs ingestion), feature allows users to provide Delphix with the location of the source database archive logs. Full backups only have to be restored during dSource creation, or during RESYNC operations when there are exceptions, for example, lost/corrupt log files. Archived logs get applied during the snapshot operation.

A full backup is required to create a dSource and prior to the snapshot operation, users can provide the archive logs to update the dSource during the snapshot operation, user-provided logs will be applied to the staging database. Applying the logs on a regular basis allows users to keep their dSource in sync with their source database. The staging database will always be in a roll forward pending state. The Plugin performs validation before applying the logs. Plugin managed Pre and Post Hooks can be used for preparing and purging the full backups and archive logs. A list of logs that are available for purging is stored in the snapshot metadata of the dSource. In case DPF is enabled, users need to place archive logs inside a folder with a name as NODE<Partition number> where <Partition number> is a four-digit number. For example, if the source environment has two partitions, then the user-provided log path will have folder names as NODE0000 and NODE0001. Both these folders will have respective archive logs. Snapshot operation will be used to apply the archive logs.

HADR and non-HADR configured dSources can be moved to this feature by performing a RESYNC operation.

Delphix supports Customer Supplied Archive Logs with Db2 version 10.5 and 11.1. To use Customer Supplied Archive Logs simply check the checkbox provided during dSource creation. Customer Supplied Archive Logs require that:

- A full backup must be accessible on the staging server, it could be a shared location but must be mounted on the staging host and visible as a file system.
- Archived logs must also be accessible on the staging server
 - The following command can be used to generate archive logs from the source database;
 - db2 archive log for database <database name>

- Instance users must have read permissions on logs so that Delphix-managed recover can apply the logs.
- Instance users must have read permissions on logs and additional write permissions on customer-provided log location if Delphix is planning to use customer-provided log location for applying the logs.
- If the user has placed the archive logs at a user-provided log location, then the logs must be valid and the first active log of the staging database must be present in the user-provided log location.
- If database configuration parameter LOGINDEXBUILD is set to OFF, do the following:
 - *update db cfg for <DB_NAME> using LOGINDEXBUILD ON*
 - Force off all connections to the database and reactivate the database
- The customer does log purging from the provided log location. Delphix only provides the list of logs that are eligible for purging.

Limitations:

- Continuous updates to the dSource are not available. The DB_SYNC operations apply new logs from the log location and keep the dSource updated with changes on the source.
- Read access to the staging database is not available.

Staging Push

With the "Staging Push" feature, the restore and rollforward operations will be managed by the user from the CLI. During dSource creation (or RESYNC) the plugin will create an empty mount point and will provide a restore and rollforward template that will help the user to restore the database. The users can restore the dSource using the steps mentioned in the readme or through any other method of their choice. Once the dSource is ready, the user can perform restore and rollforward operations from the CLI. After the restore, the staging database should remain in the rollforward pending state, and prior to the next snapshot, the user can do a rollforward to a specific Point-in-Time (by providing a timestamp to the rollforward command) against the available archive logs.

Prior to the snapshot operation, the user should validate the backups and the archive logs before applying them to the staging database. Below are the user side validations:

User side validation prior to snapshot operation

- The user must validate the backup files using **db2ckbkp** utility on the staging host.
- The user must validate the archive logs using **db2cklog** utility on the staging host.
- The authenticity of archive logs would be managed by the user.
- The user has to provide the first active log at the staging database log directory location after applying the logs.
- If there is a scenario where a user performed some load copy (non-logged transactions) operations on the source side then it's the responsibility of the user to make those non-logged transactions available on the staging database.

Backup Source Database

New Feature: Source Database with Raw DEVICE type Storage

Several users use raw device-based tablespaces for source Db2 databases. To leverage these environments with Delphix, Delphix has built a workflow using Db2s native tools that allow Delphix to discover and convert a raw device-based tablespace into an automatic storage-based tablespace during ingestion. Once the data is ingested into staging, customers will be able to provision VDBs of the automatic storage-based database.

In order to complete the linking process, the Standby dSource must have access to a full backup of the source Db2 databases on disk. This should be a compressed online Db2 backup and must be accessible to the dSource instance owner on disk. Delphix is currently not set up to accept Db2 backups taken using third-party sources such as Netbackup or TSM. All three features of data ingestion, namely HADR, Non-HADR, and Customer Supplied Archive Logs backups must also include logs. Starting with the Db2 plugin version 3.1.1, we are also supporting full online backups, which are taken using the "exclude logs" syntax. This support is only applicable to the dSources created using the Customer Supplied Archive Logs (Backup and Logs) ingestion mechanism on a single partition.

Support for named pipe type DB2 backup

Users can also provide Db2 backups of type named-pipe. Db2 plugin will check the type (type will be either named-pipe or a file) of backup and will use the restore syntax accordingly.

Example backup command: *db2 backup database <DB_NAME> online compress include logs.*

Backup types supported in DPF

The type of backups supported by Db2 plugin when used with DPF are:

- SSV backups through backup files or named pipes.
- Non-SSV backups through backup files or named pipes.

Example backup command for DB2 SSV backups: *db2_all "db2 BACKUP DATABASE<DB_NAME> ON all dbpartitionnums online to <Backup File Location> compress include logs"*

Example backup command for DB2 Non SSV backups: *db2_all "db2 backup database <DB_NAME> online to <Backup File Location> compress include logs"*

Best Practices for Taking a Backup

The following best practices can help improve backup and restore performance:

1. Compression should be enabled
2. Following parameters should be optimally configured:
 - a. Utility Heap Size (UTIL_HEAP_SZ)
 - b. No. of CPUs
 - c. No. of Table Spaces
 - d. Extent Size
 - e. Page Size
1. Parallelism & Buffer configuration may be used to improve backup performance. Parameters that should be configured are:
 - a. Parallelism
 - b. Buffer Size
 - c. No. of Buffers

More information about backup best practices is available in IBM Knowledge Center.

Procedure

1. Login to the **Delphix Management Application**.
2. On the **Databases** tab of the Environment Management screen, add a source config against the discovered staging instance.
3. Then, click **Manage** and select **Datasets**.
4. Click the **Plus (+)** icon and select **Add dSource**, you'll get a list of available source configs using which you can go for dSource creation.

5. In the **Add dSource** wizard, select the required source configuration. Consider the following when choosing an ingestion method.
 - a. Users can specify their own mount path to host the dSource dataset by specifying the path in the **Mount Base** field.
 - b. The restore and rollforward operations on the staging database can be managed by selecting the **“Use Staging Push”** checkbox. If the **“Use Staging Push”** checkbox is selected, the other checkboxes must not be selected.
 - c. If you are working with a HADR setup and want to ingest the database on a staging using HADR configurations, you need to check the **Link against source using HADR** checkbox and need to uncheck the **Customer Supplied Archive Logs** checkboxes.
 - d. To use a non-HADR method for database ingestion user need to uncheck the **Link against the source using HADR** and **Customer Supplied Archive Logs** checkboxes.
 - e. To use the **Customer Supplied Archive Logs** method for database ingestion check the **Customer Supplied Archive Logs** checkbox and uncheck the **Link against source using HADR** checkbox and enter the path in **Archive Log Path**.
 - f. To use the **Database partition feature** check **Use DPF** checkbox and uncheck the **Link against the source using HADR** checkbox. For more information on DPF please view [IBM Db2 Overview](#).

Source

Select an available environment /data source to connect to your dSource

Data Source

Filter: none ▼

☰ **R82D115A**
 new82.dcol1.delphix.com
 AppData

Data Type

AppData

Database Name (Required) *

Database name to restore from.

Mount Base

Path on the local filesystem where the dSource will be mounted by the Delphix Engine. The new path will take into effect only after executing the disable and enable operations.

Staging Push

Enabling this option will allow user to manage the staging database themselves. If this option is selected then there is no need to provide backup path and no need to select any other option.

Backup Path

Location of backup files on the Staging host. Staging host selection will be required before completing this wizard. Default backup file location is the instance home directory.

DPF (Database partitioning feature)

Enabling this option will allow creation of a dSource with DPF (Database partitioning feature) configuration

Archive Logs

Link a database, optionally providing the location of Archive Logs.

Archive Log Path

Location of Customer Supplied Archive Logs which will be applied to the staging database prior to any snapshot. Path must be present on the Staging host.

Archive Log Location is READ ONLY

Allow Delphix to copy archive logs to a temporary location, apply the logs, remove the temporary log store and then snapshot the updated staging database. This strategy may be applicable where archive logs are shared from production over NFS and the preference is to set the share to RO.

Rollforward to PIT

Rollforward a database to a particular timestamp if archive logs are provided and a valid timestamp is provided along with snapshot parameters.

Link against source using HADR

Enabling this option will configure the staging host as a HADR Standby.

HADR Primary Hostname

The host name where the primary database resides (required for HADR databases).

HADR Primary SVC

HADR Primary SVC port value (required for HADR databases).

HADR Standby SVC

HADR Standby SVC port value (required for HADR databases).

HADR Target List

HADR Target list 'host:port|host:port|host:port' values (required for HADR with multiple standby databases).

Use Plugin Hooks

Execute custom scripts when linking a dSource

Plugin Pre-Restore Hook

Execute custom scripts prior to any recovery operation.

Plugin Post-Restore Hook

Execute custom scripts following database recovery as a part of linking or resync operations.

Staging Database Configuration Settings

Customise Staging Database configuration settings. For example DATABASE_MEMORY, LOGFILSIZ.

+ Add

Cancel
Back
Next
Submit

6. The **database name** is mandatory and must be unique for a given instance. This is the name that the database was on the instance it was restored from.
7. If the user is planning to create a dSource with the "**Staging Push**" feature then the user must need to select the "**Use Staging Push**" checkbox.
8. Enter the complete **Backup Path** where the database backup file resides. If no value is entered, the default value used is the instance home directory. If there are multiple backup files for a non-DPF database on the backup path, the most current one will be used. If there are multiple sets of backup files for a DPF database on the backup path, the most recent one will be used.
9. If the user wants to roll forward a snapshot to a user-specified timestamp, then the user must select the **Rollforward to PIT** checkbox. Users can also manually do this through the snapshot option.

 The user-specified timestamp needs to be provided along with the timezone. For example, 2021-11-14-04.15.00 **UTC**. If the timezone is not entered, then the user's local timezone is considered.

 Uncheck the default snapsync policy during the dsource linking operation.

Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section.

10. Optionally, users can provide a bash script as a **Plugin** defined hook. Delphix has two Plugin managed Hooks:
 - a. **Plugin Pre-Restore Hook:** The bash script entered in this field will get executed before restoration. This hook is useful in such cases where the users want to prepare the full backup before the restore operation.
 - b. **Plugin Post-Restore Hook:** The bash script entered in this field will get executed after completion of the restore operation. This hook is useful in such cases where the user wants to purge the full backup file after the restore operation.

1. Plugin managed pre-Hook - User can copy the required full backup file from some remote location to staging host. Below is an example of such bash script (name of below script is copyBackup.sh)

```
#!/bin/bash
```

```
# Copying backup file from a remote host remote.delphix.com
scp auto1051@remote.delphix.com:/backuplocation/
R74D105A.0.auto1051.DBPART000.20180510073639.001 /db2backup/
```

The user can provide the above script in Plugin managed pre-Hook field as "sh copyBackup.sh"

2. Plugin managed post-Hook - User can purge the full backup file from the staging host. Below is an example of such bash script (name of below script is purgeBackup.sh)

```
#!/bin/bash
```

```
# Purging full backup file from staging host for saving the staging storage
rm -f /db2backup/R74D105A.0.auto1051.DBPART000.20180510073639.001
```

The user can provide the above script in Plugin managed post-Hook field as "sh purgeBackup.sh"

11. If the dSource is to use **HADR** please enter the following fields. For more information about HADR please view [Linking a dSource from a Db2 Database: An Overview](#).
 - a. Enter a fully qualified HADR Primary Hostname. This is a required field for HADR and must match the value set for HADR_LOCAL_HOST on the master.
 - b. Enter the port or /etc/services name for the HADR Primary SVC. This is a required field for HADR and uses the value set for HADR_LOCAL_SVC on the master.
 - c. Enter the port or /etc/services name for the HADR Standby SVC. This is a required field for HADR and uses the value set for HADR_REMOTE_SVC on the master.
12. If you are using **Customer Supplied Archive Logs** for dSource creation complete the following fields:
 - a. Check the Customer Supplied Archive logs checkbox
 - b. Provide the log location of the archive logs. In this location, users can load copy files by appending schema and table name to it so that plugin can process those load copy files during log apply process. Archive logs will be used for applying the logs during the snapshot operation. This location must be a part of the staging server and this location must be different than full backup location.

The screenshot shows the 'Add dSource' configuration window. On the left, a sidebar lists steps: Source (selected), dSource Configuration, Data Management, Policies, Hooks, and Summary. The main content area is divided into two columns. The right column contains the following options and fields:

- Customer Supplied Archive Logs
Link a database, optionally providing the location of Archive Logs.
- Archive Log Path
[Text input field]
- Location of Customer Supplied Archive Logs which will be applied to the staging database prior to any snapshot. Path must be present on the Staging host.
- Archive Log Location is READ ONLY
Allow Delphix to copy archive logs to a temporary location, apply the logs, remove the temporary log store and then snapshot the updated staging database. This strategy may be applicable where archive logs are shared from production over NFS and the preference is to set the share to RO.
- Link against source using HADR
Enabling this option will configure the staging host as a HADR Standby.
- HADR Primary Hostname
[Text input field]
- The host name where the primary database resides (required for HADR databases).
- HADR Primary SVC
[Text input field]
- HADR Primary SVC port value (required for HADR databases).
- HADR Standby SVC
[Text input field]
- HADR Standby SVC port value (required for HADR databases).
- HADR Target List
[Text input field]

- c. The **Archive Log Location is READ ONLY** (default) option indicates whether Delphix should copy the log files to the dSource’s filesystem prior to the log apply process. If this is not selected the staging database will read the log files directly from the user-provided log location and perform the log apply operation.

⚠ Special considerations for Archive Log Location is READ ONLY

If this operation is performed on the source database and user want the Plugin to process the load files place the load copy files in your location and append the schema and table names (for example: <DB Name>.<Number>.<Instance>.DBPART<partNumber>.<Timestamp>.001.<schema name>.<table name>) to the load copy file names
If you do not want to process load copy files via the Plugin place the load copy files in a similar structure as it is on the source or set a DB2 registry parameter DB2LOADREC.

⚠ Special consideration during dSource upgrade from older version plugin to Customer Supplied Archive Logs feature enabled plugin

Suppose the user is having a dSource with an older version of the plugin and now wants to upgrade that dSource to **Customer Supplied Archive Logs** feature then the user needs to perform the following steps :

1. Go to the configuration tab of dSource and then click on the custom sub-tab.
2. Click on the edit icon.
3. Check the **Customer Supplied Archive Logs** and **Archive Log Location is READ ONLY** checkbox.
4. Click on submit and then go for RESYNC operation to upgrade the dSource.

13. Click **Next**.
14. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
15. Click **Next** and set the **Staging Environment** to be the same as the dSource host. Select the **Staging Environment User** to be the same as the instance owner of the dSource instance.

✔ **Changing the Environment User**

If you need to change or add an environment user for the dSource instance, see [Managing Db2 Users and Instance](#).

⚠ Ensure that the Staging environment must be the environment that you have chosen to create the source config.

Select the **User** from the drop-down and set the **SNAPSHOT PARAMETERS** for the dSource snapshot.

- a. Click on the **Timestamp** input field to enter the timestamp at which the snapshot will be taken. The user-specified timestamp needs to be provided along with the timezone in YYYY-MM-DD-HH.MM.SS format. For example, 2022-03-12-04.15.00**UTC**. If the timezone is not entered, then the user's local timezone is considered.

⚠ Providing the timestamp will only work if the user has checked the **Rollforward to PIT** option in step 9 above.

- b. Select **Resynchronize dSource** to resynchronize the dSource. This will force a non-incremental load of data from the source similar to when the dSource was created. This action avoids duplicating storage requirements and maintains a Timeflow history.
16. The Delphix Engine will initiate two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking Active Jobs in the top menu bar, or by selecting System > Event Viewer. When the jobs have completed successfully, the database icon will change to a dSource icon on the Environments > Host > Databases screen, and the dSource will also appear in the list of Datasets under its assigned group. Then, click **Next** and you'll get the Policies section. Set the desired **Snapsync Policy** for the dSource. For more information on policies see Advanced Data Management Settings for Db2 dSources.
 17. Click **Next** and specify any desired pre- and post-scripts. For details on pre- and post-scripts, refer to Customizing Db2 Management with Hook Operations.
 18. Click **Next**. Review the dSource Configuration and Data Management information in the Summary section.
 19. Click **Submit**.

ℹ **The dSource Configuration Screen**

After you have created a dSource, the **dSource Configuration tab** allows you to view information about it and make modifications to its policies and permissions. In the **Datasets** panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the **Source files, Data Management** configuration, and **Hook Operations**.

⚠ **Purging of archive logs after snapshot operation**

Logs can be purged from your log location after the snapshot operation. The Plugin will preserve the list of logs that are eligible for purging in the snapshot's metadata. The snapshot metadata can be accessed via the Delphix CLI or API.

Backup Source Database

Source Database with Raw DEVICE type Storage

Several users use raw device-based tablespaces for source DB2 databases. To leverage these environments with Delphix, Delphix has built a workflow using Db2s native tools that allow Delphix to discover and convert a raw device-based tablespace into an automatic storage-based tablespace during ingestion. Once the data is ingested into staging, customers will be able to provision VDBs of the automatic storage-based database.

In order to complete the linking process, the Standby dSource must have access to a full backup of the source Db2 databases on disk. This should be a compressed online DB2 backup and must be accessible to the dSource instance owner on disk. Delphix is currently not setup to accept DB2 backups taken using third-party sources such as Netbackup or TSM. All three features of data ingestion, namely HADR, Non-HADR, and Customer Supplied Archive Logs backups must also include logs. Starting with Db2 plugin version 3.1.1, we are also supporting full online backups, which are taken using the "exclude logs" syntax. This support is only applicable to the dSources created using the Customer Supplied Archive Logs (Backup and Logs) ingestion mechanism on a single partition.

Example backup command: `db2 backup database <DB_NAME> online compress include logs`

Best Practices for Taking a Backup

The following best practices can help improve backup and restore performance:

1. Compression should be enabled
2. Following parameters should be optimally configured:
 - a. Utility Heap Size (UTIL_HEAP_SZ)
 - b. No. of CPUs
 - c. No. of Table Spaces
 - d. Extent Size
 - e. Page Size
1. Parallelism & Buffer configuration may be used to improve backup performance. Parameters that should be configured are :
 - a. Parallelism
 - b. Buffer Size
 - c. No. of Buffers

More information about backup best practices is available in [IBM Knowledge Center](#).

Procedure

1. Login to the Delphix Management Application using Delphix Engine credentials or as the owner of the database from which you want to provision the dSource.
2. On the **Databases** tab of Environment Management screen, add a source config against discovered staging instance.
3. Then, click **Manage**.
4. Select **Datasets**.
5. Click the Plus (+) icon and select **Add dSource**, you'll get a list of available source configs using which you can go for dsource creation.
6. In the **Add dSource** wizard, select the required source configuration.
7. If you are working with an HADR setup, please leave the HADR checkbox checked.
8. The **database name** is mandatory and must be unique for a given instance. This is the name that the database was on the instance it was restored from.

9. Enter the complete **Backup Path** where the database backup file resides. If no value is entered, the default value used is the instance home directory. If there are multiple backup files for a database on the backup path, the most current one will be used.
10. Enter the **Log Archive Method1** you wish to use for the database. If no value is entered, the default value used is `DISK:/mountpoint/dbname/arch`.
11. Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section.
12. If the dSource is to use HADR please enter the following fields. If it will not use HADR skip ahead to step 13. For more information about HADR please view [Linking a dSource from a Db2 Database: An Overview](#).
 - a. Enter a fully qualified HADR Primary Hostname. This is a required field for HADR and must match the value set for HADR_LOCAL_HOST on the master.
 - b. Enter the port or /etc/services name for the HADR Primary SVC. This is a required field for HADR and uses the value set for HADR_LOCAL_SVC on the master.
 - c. Enter the port or /etc/services name for the HADR Standby SVC. This is a required field for HADR and uses the value set for HADR_REMOTE_SVC on the master.
13. Click **Next**.
14. Select a **dSource Name** and **Database Group** for the dSource.
15. Click **Next**.
You will get Data Management section where you need to specify staging environment and user which will be used for dsource creation.
16. Set the **Staging Environment** to be the same as the dSource host.
17. Select the **Staging Environment User** to be the same as the instance owner of the dSource instance.

 **Changing the Environment User**

If you need to change or add an environment user for the dSource instance, see [Managing Db2 Users and Instance Owners](#).

18. Then, click **Next** and you'll get the Policies section. Set the desired **Snapsync Policy** for the dSource.
19. Click **Next**.
20. Specify any desired pre- and post-scripts. For details on pre- and post-scripts.
21. Click **Next**.
22. Review the dSource Configuration and Data Management information in the summary section.
23. Click **Submit**.

The Delphix Engine will initiate two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking Active Jobs in the top menu bar, or by selecting System > Event Viewer. When the jobs have completed successfully, the database icon will change to a dSource icon on the Environments > Host > Databases screen, and the dSource will also appear in the list of Datasets under its assigned group.

 **The dSource Configuration Screen**

After you have created a dSource, the **dSource Configuration tab** allows you to view information about it and make modifications to its policies and permissions. In the **Datasets** panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the **Source files, Data Management** configuration, and **Hook Operations**.

Provisioning a Virtual Database (VDB)

Prerequisites

- You will need to have linked a dSource from a staging instance, as described in [Linking a Db2 dSource](#), or have created a VDB from which you want to provision another VDB

- You should have set up the Db2 target environment with necessary requirements as described in [Requirements for Db2 Hosts and Databases](#)
- Make sure you have the required Instance Owner permissions on the target instance and environment as described in [Managing Db2 Users and Instance Owners](#)
- The method for [Database Permissions for Provisioned Db2 VDBs](#) is decided before the provisioning
- [Optional] Once you have restored the databases into your target environment, If the version of IBM® Db2® in your source environment is at a different fix pack level than the version of Db2 in your target environment, then you need to update the 11.1 version databases with db2updv111 command and 11.5 databases with db2updv115 command utility.

Run the below command as a Pre-snapshot hook:

```
db2updv111 -d <DBNAME> -a
db2updv115 -d <DBNAME> -a
```

The optional pre-requisite is applicable only when the target environment is using a different Db2 fix pack level than the source. If your source environment is using version 11.1.x and your target is using version 11.5.x, you can neglect the above-mentioned point.

Consider the following scenarios for the optional prerequisite mentioned above:

- if your source environment uses Db2 version 11.5.5 and your target environment uses Db2 11.5.8 or a later fix pack, you need to update the databases in the target environment.
- If your databases were created or upgraded to Db2 Version 11.1 GA and you applied Db2 Version 11.1 Fix Pack 1 or later, then running the above command will automatically apply all updates required from Version 11.1 GA up to and including the fix pack level that you are installing.



- To know how to update databases, see [db2updv115 - Update database to Version 11.5 mod pack command](#).
- Certain database or instance operations are restricted while an online fix pack update is in progress.
- See [db2updv111 - Update database to Version 11.1 fix pack command](#), which mentions it is mandatory to run db2updv111 on database if one is coming from higher version of fix pack to Lower OR vice versa.
- For more information, you can also refer: [Update the databases for a Db2 11.1.x fix pack](#)



You can take a new snapshot of the dSource by clicking the **Camera** icon on the dSource card. Once the snapshot is complete you can provision a new VDB from it.



Please note while creating a VDB on an instance

By default, DB2 Plugin will not allow provisioning a VDB on an instance which contains a database that has an identical name to the source database.

For e.g.;

1. When we are provisioning a VDB from a VDB then source VDB will be treated as the source database.
2. When we are provisioning a VDB from a dSource then a staging database will be treated as the source database.

If the user wants to provision a VDB on an instance which contains a source database, then the user needs to follow the below steps :

1. Check if any config file with name db2_plugin.conf at location **<Toolkit directory>/Delphix_COMMON/plugin/<Identifier>/**. If it does not exist, then the user should create a file with name db2_plugin.conf at location **<Toolkit directory>/Delphix_COMMON/plugin/<Identifier>/**.

The instance user must have read permission to this file. Once the file is created, user should proceed to step 2. If the file already exists at this location with the required permission, the user should directly proceed to step 2.

2. [Add the below line in this config file.](#)

1. allowSourceDbOnSameInstance=**True**

- [Then create/Refresh the VDB again.](#)

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource**.
5. Select a **snapshot** from which you want to provision.
6. Click **Provision VDB** icon to open Provision VDB wizard.
7. Select a target environment from the left pane.
8. Select an **Installation** to use from the dropdown list of available DB2 instances on that environment.
9. Set the **Environment User** to be the Instance Owner. Note: The picking of instance owner is only possible if you have multiple environment users set on that host.
10. Provide VDB Name as database name as parameter.
11. Optionally, set the database configuration parameters for the VDB.
12. Click **Next**.
13. Select a **Target Group** for the VDB.
Click the green **Plus** icon to add a new group, if necessary.
14. Select a **Snapshot Policy** for the VDB.
15. Click **Next**.
16. Specify any desired hook operations.
17. Click **Next**.
18. Review the **Provisioning Configuration** and **Data Management** information.
19. Click **Submit**.

When provisioning starts, you can review the progress of the job in the **Databases** panel, or in the **Job History** panel of the **Dashboard**. When provisioning is complete, the VDB will be included in the group you designated and listed in the **Databases** panel. If you select the VDB in the Databases panel and click the **Open** icon, you can view its card, which contains information about the database and its Data Management settings.

Once the VDB provisioning has successfully completed, if the source and target instance ids are not the same, you may want to grant secadm and dbadm on the database to the target instance id. Refer to [Database Permissions for Provisioned Db2 VDBs](#) for more information.

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.
2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision new VDBs from that masked VDB to quickly provide safe data to development and QA teams.

Related Topics

- [Linking Data Sources and Syncing Data with Db2](#)
- [Provisioning and Managing VDBs from Db2 dSources](#)

Db2 Support and Requirements

In this section, we will describe all the different Delphix and Db2 components that are part of the Delphix Virtualization solution for IBM Db2 database.

- [Db2 Compatibility Matrix](#)
- [Network and Connectivity Requirements for Db2 Environments](#)
- [Requirements for Db2 Hosts and Databases](#)
- [Sudo Privilege Requirements for Db2 Environments](#)
- [Sudo File Configuration Examples for Db2 Environments](#)

Db2 Compatibility Matrix

This topic describes the Db2 (DBMS) versions that are supported by Delphix, as well as the compatible operating systems (OS), for use on target and source environments.

Supported Db2 Database Editions

1. Enterprise Server Edition
2. Advanced Enterprise Server Edition
3. Advanced Edition
4. Standard Edition
5. For the supported Db2 versions, Delphix supports the corresponding Db2 Developer edition with Db2 plugin support

-  • Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Db2 versions below 10.5 i.e (10.1, 9.8) are NOT supported by the plugin.
- Db2 supports only 64-bit OS

Key:

Color	Supported?
	Yes
	No

Red Hat Enterprise Linux (RHEL)

-  Support for RHEL 8.x requires installing the `libncurses.5` library onto the host. Please reference [KBA 5622](#) for actionable steps.

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
RHEL 6.9			
RHEL 7.0			
RHEL 7.1			

RHEL 7.2	✓	✓	✗
RHEL 7.3	✓	✓	✗
RHEL 7.4	✓	✓	✗
RHEL 7.5	✓	✓	✗
RHEL 7.6	✓	✓	✗
RHEL 7.7	✗	✓	✗
RHEL 7.8	✗	✓	✓ 6.0.5+
RHEL 8.2	✗	✗	✓

Advanced Interactive eXecutive (AIX)

Db2 plugin support for below mentioned OS and Db2 versions

Supported OS Version	Supported DBMS Version		
	10.5	11.1	11.5
AIX 7.1	✓	✓	✓
AIX 7.2	✓	✓	✓

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

 The plugin versions starting 2.1.0 till 2.7.3 are on extended support. Please contact Delphix Support for any queries.

Plugin Version

Delphix Engine	3.0.0	3.0.1	3.0.2	3.1.0	3.1.1	4.0.0	4.0.1	4.1.0	4.1.1	4.1.2	4.1.3
6.0.0.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.1.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.2.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.3.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.4.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.5.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.6.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.7.x	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Related Topics

- [Db2 Support and Requirements](#)

Network and Connectivity Requirements for Db2 Environments

Overview

This topic outlines the network and connectivity requirements for the Delphix Engine and Db2 standby and target environments.

Port Allocations Specific to DB2

The Delphix Engine makes use of the following network ports for Db2 standby and target:

Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP/UDP	111	Remote Procedure Call (RPC) port mapper used for NFSv3 mounts Note: RPC calls in NFSv3 use additional fixed ports for supporting services (lockd, mountd and statd) seen below.
TCP	1110	NFS Server daemon status and NFS server daemon keep-alive (client info)
TCP	2049	NFS Server daemon from vFiles to the Delphix Engine
TCP	54043	NFSv3 mount daemon
TCP	54044	NFSv3 stat daemon (lock state notification service)
TCP	54045	NFSv3 lock daemon/manager
UDP	33434 - 33464	Traceroute from standby and target hosts to the Delphix Engine (optional)

Outbound from a Standby or Target Environment Port Allocation

Protocol	Port Numbers	Use
TCP	873	Rsync connections used during V2P.
TCP	8415	DSP connections used for monitoring and script management. Typically DSP runs on port 8415.

Inbound to a Standby or Target Environment Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to the target environment

HADR Service Ports

The HADR ports set for HADR_LOCAL_SVC and HADR_REMOTE_SVC on the Db2 Master and Standby hosts. The specific ports used at the customers' discretion and need to be specified during the linking process. It is highly recommended that these ports also be defined in the /etc/services file to ensure that they are only used by Db2 for the specified databases.

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI

Protocol	Port Number	Use
UDP	161	Messages from an SNMP Manager to the Delphix Engine
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

AppData Port Requirements The use of AppData requires the following ports/protocols.

Two important notes about these specifications:

1. The next release of the Delphix Engine will significantly augment the port/protocol utilization of AppData. The upcoming-only requirements have been marked with a *.
2. AppData V2P uses RSYNC to export to the target. RSYNC between the target and Delphix Engine is not required for general virtualization usage. The V2P-only requirements have been marked with a ^.

From Source to Delphix Engine	From Delphix Engine to Source	From Target to Delphix Engine	From Delphix Engine to Target
RSYNC (TCP Port 873)	RSYNC (TCP Port 873)	DSP (Default TCP Port 8415)	DSP (Default TCP Port 8415)
DSP (Default TCP Port 8415)	SSH (TCP Port 22)	NFS	SSH (TCP Port 22)
*NFS	DSP (Default TCP Port 8415)	^RSYNC (TCP Port 873)	^RSYNC (TCP Port 873)

Related Topics

- [Db2 Support and Requirements](#)

Requirements for Db2 Hosts and Databases

Source Db2 hosts are servers that have Db2 binaries installed and have Db2 instances created on them. The hosts that contain the data that we wish to ingest are referred to as the source environment. Hosts with empty instances (no dbs in instance) are used as either staging or target hosts. This topic describes the requirements for creating connections between the Delphix Engine and Db2 hosts and instances.

Requirements for Db2 Source Hosts and Instances

Each Db2 Source host must have IBM Db2 installed and at least one instance created on the machine. Depending on the way you intend to ingest data into the Staging server, you will also need to ensure that the following requirements are met:

- For HADR Ingestion: HADR settings for each database to be used with the standby server should be preset before the linking process begins as described in [Linking a Db2 dSource](#)
- For Backup and Log Ingestion, no additional requirements are needed at this point.

Requirements for Db2 Staging and Target Hosts and Instance

- The staging environment that the Delphix Engine uses must have access to an existing full backup of the source database on disk to create the first full copy. Delphix recommends using compressed backups as that will reduce storage needs and speed up ingest. Once this first full backup is available, subsequent refreshes can be accomplished via HADR or Archive Logs.
- The staging and target Db2 instances that you wish to use must already exist on the host. We can use the same instance for the dSource and VDB creation and can also have multiple VDBs on the same instance.
- The available instances on each host can be verified by going to the databases tab for the environment in question.
- Instance level configuration values such as the bufferpool value will need to be managed by the customer independent of Delphix.
- The instances used for staging and target environments must be compatible with the source Db2 instance.

Requirements for Db2 DPF Staging and Target Hosts and Instances

i It is highly recommended that the Database Partitioning Feature (DPF) for Db2 staging and target should be configured on separate hosts due to reason that the DPF environment consumes a lot of resources.

- The staging and target hosts should be configured with the same number of logical partition nodes as the source.
- The logical portion configuration should be added in \$HOME/sqllib/db2nodes.cfg file where \$HOME is the home directory of the instance user.
- The hostname to be used should be a fully qualified domain name.
- db2 nodes.cfg file configured for 4 logical partitions should look like:

```
0 <FQDN hostname> 0
1 <FQDN hostname> 1
2 <FQDN hostname> 2
3 <FQDN hostname> 3
```

- Each logical partition should have a corresponding port entry in /etc/services of hosts. In the example below, db2inst1 is the instance the user and port number from 60000 to 60003 are reserved for inter-partition communication.

```
DB2_db2inst1 60000/tcp
DB2_db2inst1_1 60001/tcp
DB2_db2inst1_2 60002/tcp
DB2_db2inst1_END 60003/tcp
```

- Registry variables DB2RSHCMD and DB2COMM should be updated with values as shown below:

```
SSHCMD=$(which ssh)
db2set DB2RSHCMD="$SSHCMD"
db2set DB2COMM=TCPIP
```

- After the environment is configured on staging and target, execute **db2_all echo hello** on both staging and target servers to add host key in the known_hosts file. This will also help to check for any ssh issues on these servers.
- The backup files (or log directory if using **Customer Supplied Archive Logs**) should be consistent with the partitions configured on staging.

Additional Environment Requirements

- There must be an operating system user (delphix_os) with these privileges:
 - Ability to login to the staging and target environment via SSH
 - Ability to run mount, umount, mkdir, and rmdir as a super-user. If the staging and target host is an AIX system, permission to run the nfsd command as a super-user. See [Sudo Privilege Requirements for Db2 Environments](#) for further explanation of the commands and [Sudo File Configuration Examples for Db2 Environments](#) for examples of the /etc/sudoers file on different [Outdent](#) operating systems.
- There must be a directory on the staging and target environment where you can install the Delphix Engine Plugin. For example, /var/opt/delphix/toolkit .
 - The delphix_os user must own the toolkit path directory (__/var/opt/delphix/toolkit).
 - If the delphix_os user and instance users (responsible for the Delphix Engine operations such as linking and provisioning) are sharing any common group, then the toolkit path directory (__/var/opt/delphix/toolkit) must have permissions rwxrwx-- (0770), but you can also use more permissive settings. This allows instance users who are part of the same group as the delphix_os user's group to be able to create directories inside the toolkit path directory.
 - If the delphix_os user and instance users (responsible for the Delphix Engine operations such as linking and provisioning) are not sharing any common group, then the toolkit path directory must have -rwxrwxrwx (0777) permissions.
 - The directory should have 1.5GB of available storage: 400MB for the Plugin and 400MB for the set of logs generated by each Db2 instance that runs on the host
 - In Db2 Plugin, plugin directory space will be used as the base location for the mount point
 - When multiple engines access the same host environment, create a separate Db2 plugin directory for each engine that will discover this environment.

- i** For Db2 plugin version 4.0.0 and onwards, follow the steps below to change the toolkit directory.
- Disable all the datasets related to that environment
 - The instance of the dataset must have the DIR_CACHE configuration parameter set to No, as shown below.
db2 update dbm cfg using DIR_CACHE NO
 - Upload the new vSDK-based plugin. The plugin has upgrade logic implemented in it and will upgrade the existing Lua-based plugin.

- iv. Refresh the environment to allow the plugin to create plugin paths and files under the scratch path.
- v. Enable all the datasets and complete the migration process.

- The Delphix Engine must be able to initiate an SSH connection to the staging and target environment
- NFS client services must be running on the staging and target environment
- Staging and Target host instances must have the same Db2 instance code page as the production instance. The instance owner must have a compatible value defined in LC_ALL or LC_CTYPE for the Db2 code page. Preferably derived from the value of LC_ALL, LC_CTYPE, or LANG of the instance owner on the production host.
- The primary environment user should share its primary group with the instance user. For example, if the delphix_os is the primary environment user which is used for environment addition and its primary group is also delphix_os then instance users (responsible for the Delphix Engine operations such as linking and provisioning) should share group delphix_os as their secondary group.

Custom Configurations in Db2 Plugin

DB2 Plugin v3.x

The advanced config file is named **advanceConfFlag.conf** and is created under the **<Toolkit directory>** during discovery or environment refresh operation if the file does not pre-exist already. This file will have the following configurable parameters:

- **Common Group (*notCommonGroupFlag*):**
 - This parameter sets the permission of "mnts", "logs" and "code" directories under the **<Toolkit dir>/DB2** location.
 - Set ***notCommonGroupFlag*** to *false*, if the primary group of the primary environment user (user which is used to do discovery) is shared with the instance user.
 - Set ***notCommonGroupFlag*** to *true*, if the primary environment user and the instance user do not share any user group.
 - Once the necessary changes are made, refresh the environment in order to apply the changes.

By default, the ***notCommonGroupFlag*** parameter is commented out in the file. This means that the plugin implicitly assumes its value as false. The valid values for this parameter are true or false.

- **Allow Source Database on Same Instance (*allowSourceDBonSameInst*)**

By default, the DB2 Plugin restricts the provisioning of a VDB on a DB2 instance which contains a database with the same name as the VDB's source database

For example:

- When provisioning a VDB from a VDB then the latter is treated as the source database.
- When provisioning a VDB from a dSource then the staging database representing the dSource is treated as the source database.

Set ***allowSourceDBonSameInst*** to true, if the user wants to provision a VDB on an instance which contains a database with the same name as the VDB's source database. The valid values for this parameter are true or false.

- **Restore Pipeline limit (*restorePipelineLimit*):**

Db2 DPF Plugin performs restoration of all non-catalog partitions in parallel.

The parameter ***restorePipelineLimit*** allows the user to configure the number of partitions that could be restored in parallel by the plugin. When the new **advanceConfFlag.conf** file is created during environment refresh or discovery operation, the default value of

restorePipelineLimit is set to 10. The user can modify the value of this parameter to a desired value. The valid value for this field is any positive integer.

A default config file created during discovery or environment refresh operation:

```
# If the user is not sharing a common group between primary environment user and
instance user.

# Then the user needs to set parameter notCommonGroupFlag as true.

# notCommonGroupFlag=true

# During provision operation check if instance contains a database name which is
identical to the source DB name used for provisioning operation.

# By default plugin will never allow creating a VDB on the instance where source
database (or other database which is identical in name with source DB) already
exists.

# If the user still wants to create a VDB on the instance which contains a database
name which is identical to the source DB name used for provisioning operation,

# then the user needs to set parameter allowSourceDBonSameInst as true.

# allowSourceDBonSameInst=true

# Limit for parallel restores

restorePipelineLimit=10
```

DB2 Plugin v4.x and onwards

DB2 plugin config file is named **db2_plugin.conf** and is created under the **<Toolkit directory>/<Delphix_COMMON>/plugin/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/** directory during discovery or environment refresh operation if the file does not already exist. This file has two sections represented by the headers - `plugin_custom_parameters` and `plugin_logging_parameters`.

 It is important for the users to maintain two section headers in the `db2_plugin.conf` file. If any of the sections get deleted, the plugin will pick the default value for the parameters under the deleted section.

Plugin Custom Parameters

Common Group (`usersHaveCommonGroup`)

- This parameter sets the permission for the "mnts" directory located under the **<Toolkit dir>/DB2** directory and for the "logs" directory located under the **<Toolkit directory>/<Delphix_COMMON>/plugin/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/** directory.
- Set the `usersHaveCommonGroup` parameter as **True**, if the primary group of the primary environment user (user which is used to do discovery) is shared with the instance user.
- Set the `usersHaveCommonGroup` parameter as **False**, if the primary environment user and the instance user do not share any user group.

- Once the necessary changes are made, refresh the environment in order to apply the changes.

By default, the **usersHaveCommonGroup** parameter is set to **True**. The valid values for this parameter are true or false.

Allow Source Database on Same Instance (*allowSourceDbOnSameInstance*)

By default, the DB2 Plugin restricts the provisioning of a VDB on a DB2 instance which contains a database with the same name as the VDB's source database

For example:

- When provisioning a VDB from a VDB then the latter is treated as the source database.
- When provisioning a VDB from a dSource then the staging database representing the dSource is treated as the source database.

If the user wants to provision a VDB on an instance which contains a database with the same name as the VDB's source database, then the parameter **allowSourceDbOnSameInstance** must be set to true. The valid values for this parameter are true or false.

Restore Pipeline limit (*restorePipelineLimit*)

Db2 DPF Plugin performs restoration of all non-catalog partitions in parallel.

The parameter **restorePipelineLimit** allows the user to configure the number of partitions that could be restored in parallel by the plugin. When the new **advanceConfFlag.conf** file is created during environment refresh or discovery operation, the default value of **restorePipelineLimit** is set to 10. The user can modify the value of this parameter to the desired value. The valid value for this field is any positive integer.

Plugin Logging Parameters

This feature provides three options to configure and control enhanced logging by the plugin:

1. To set debug log levels.
2. To set active log file size.
3. To set log retention levels.

Delphix users can modify the above parameters defined under the `plugin_logging_parameters` section of this file.

Setting Log Levels (level)

Delphix users can set the logging level of the plugin-generated logs by setting the parameter **level** under the `plugin_logging_parameters` header. Higher logging levels will help to expedite debugging issues. There are two levels of logging which are: Info and Debug. The configured logging level will apply to all the objects (both dSources and VDBs) present on a staging/target host. For example; if we have 4 dSources on a host associated with a Delphix Engine then the log level will apply to all the dSources.

Log level description:

- level=INFO: This level will print only informational logs. This is the default log level.
- level=DEBUG: This level will print informational logs and debug statements.

Setting Maximum Size of Active Log File (logFileSize)

Delphix users can set the maximum size of the active log file by defining the value of the `logFileSize` parameter under the `plugin_logging_parameters` header of the `db2_plugin.conf` file. Once this maximum size limit is reached, the plugin will rotate the log as per the retention property described below. This parameter only takes in a positive integer value. The minimum value of this parameter is 1 MB and the maximum value is 10 MB. The default value is 1 MB.

Setting Log Retention Levels (retention)

Delphix users can set a retention level for plugin-generated logs for each dataset (<Db name>.diag.log) using the parameter **retention** under the plugin_logging_parameters header. As per this parameter, the log files are moved (archived), renamed or deleted once they reach the value set in the logFileSize parameter. New incoming log data is directed into a new fresh file (at the same location).

By default, the value for this parameter will be set to a minimum value of 2. The user can change this value and set its value within the range 2 and 50. For example; if **retention** is set to 4, the plugin will have the following log files: <DB Name>.diag.log, <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB Name>.diag.log.3, <DB Name>.diag.log.4.

- File <DB Name>.diag.log is the active log file
- File <DB Name>.diag.log.1 is the most recent archive log file
- File <DB Name>.diag.log.4 is the oldest one

 The DB2 plugin will ignore values for a parameter if the values do not fit the description above and will continue the operation with the default set of values for that particular parameter.

db2_plugin.conf

```
#
# The parameters must not be left padded with spaces. The plugin will otherwise not
# accept the specified parameter
# values and will use the default values.
#

[plugin_custom_parameters]
#
# If the user has not shared a common group between primary environment user and
# instance user.
# Then the user needs to set parameter usersHaveCommonGroup as False.
#

usersHaveCommonGroup=True

#
# During provision operations the plugin checks if the target instance contains a
# database name which is identical to
# the source DB name from which the VDB is to be created. By default the plugin will
# never allow creating a VDB on the
# instance where the source database already exists, or on any instance that has a
# database with the same name. If the
# users still want to create a VDB on an instance which contains a database with the
# same name as the source DB then
# they need to set parameter allowSourceDbOnSameInstance as True.
#

allowSourceDbOnSameInstance=False

#
# If the users want to control the number of parallel restores then they can tune the
# parameter restorePipelineLimit.
# By default the value of this parameter would be 10.
```

```

#

restorePipelineLimit=10

[plugin_logging_parameters]
#
# This flag will set the debug level of plugin logs on the remote server (both
staging and target). There are two levels:
# Info
# Debug
# The above are the only valid values that can be assigned. If any other value is
assigned, the plugin will set the
# default level as Info. the string is case insensitive so info, Info and INFO are
acceptable.
#

level=INFO

#
# This parameter will set the maximum size of the active log file in MB. Once this
limit is reached, the plugin will
# rotate the log as per the retention property defined below. The minimum value of
this parameter is 1 MB and maximum
# value is 10 MB. This parameter only takes in a positive integer value. The default
value is 1 MB
# logFileSize=<positive integer>
#

logFileSize=1

#
# Whenever the size of <DB Name>.diag.log file exceeds the value provided by
logFileSize parameter then the plugin will
# rename the active log file to <DB Name>.diag.log.<number> and a new log file with
name <DB Name>.diag.log will be
# generated. For example, if LogRetention is set to 4, the plugin will have the
following log files: <DB Name>.diag.log,
# <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB Name>.diag.log.3, <DB
Name>.diag.log.4.
# File <DB Name>.diag.log.4 will be the oldest one.
# File <DB Name>.diag.log.1 will be the most recent archive log file.
# File <DB Name>.diag.log will be the active log file.
# The minimum and default value for retention flag is 2.
# retention=<positive integer>
#

retention=2

```

Instance User Requirements

- The instance owner of each instance you wish to use within staging or a target host must be added as an environment user within the Delphix engine. See [Managing Db2 Users and Instance Owners](#).

- For HADR synced dSources the staging instance owner must be able to "read" the ingested database contents as Delphix will check the validity of the database by querying tables before each dSource snapshot.

Database Container Requirements

- All Db2 database container types are fully supported with the exception of Db2 raw containers.
- It is recommended not to use special characters for source database container paths as this may create problems while parsing such container paths and may also result in failures during dSource and VDB creation.

 If a container is added or deleted, the dSource will have to be resynced.

Related Topics

- [Db2 Support and Requirements](#)

Sudo Privilege Requirements for Db2 Environments

This topic describes the rationale behind specific `sudo` privilege requirements for virtualizing Db2 Databases.

Privilege	Sources	Targets and Staging	Rationale
<code>mkdir/rmdir</code>	Not Required	Required	Delphix dynamically makes and removes directories under the provisioning directory during VDB operations. This privilege is optional, provided the provisioning directory permissions allow the delphix os user to make and remove directories.
<code>mount/umount</code>	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for superuser.

 It is required to specify the NOPASSWD qualifier within the "sudo" configuration file, as shown here: [Sudo File Configuration Examples for Db2 Environments](#). This ensures that the "sudo" command does not demand the entry of a password, even for the "display permissions" (i.e. "sudo -l") command. Delphix issues "sudo -l" in some scripts to detect if the operating system user has the correct sudo privileges. If it is unable to execute this command, some actions may fail and Delphix will raise an alert suggesting it does not have the correct sudo permissions. Restricting the execution of "sudo -l" by setting "listpw=always" in the "/etc/sudoers" file when the Delphix operating system user is configured to use public key authentication will cause the Delphix operating system user to be prompted for a password which will fail certain Delphix actions. Use a less restrictive setting for listpw than "always" when the Delphix operating system user is using public-key authentication.

Related Topics

- [Db2 Support and Requirements](#)

Sudo File Configuration Examples for Db2 Environments

Configuring `sudo` Access on AIX for Db2 Source and Target Environments

Sudo access to `ps` on the AIX operating system is required for the detection of listeners with non-standard configurations on both source and target environments. Super-user access level is needed to determine the `TNS_ADMIN` environment variable of the user running the listener (typically **db2**, the installation owner). From `TNS_ADMIN`, the Delphix OS user **delphix_os** can derive connection parameters.

Example: AIX /etc/sudoers entries for a Delphix Source

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/bin/ps
```

In addition to sudo access to the `mount`, `umount`, and `ps` commands on AIX target hosts, Delphix also requires `sudo` access to `nfso`. This is required on target hosts for the Delphix Engine to monitor the NFS read-write sizes configured on the AIX system. Super-user access level is needed to run the `nfso` command.

Example: AIX /etc/sudoers File for a Delphix Target

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
/usr/sbin/mount, \
/usr/sbin/umount, \
/usr/sbin/nfso, \
/usr/bin/ps
```

Configuring `sudo` Access on Linux for Db2 Source and Target Environments

On a Linux target, sudo access to `mount`, `umount`, `mkdir`, and `rmdir` is required.

Example: Linux /etc/sudoers file for a Delphix Target for DB2

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
/bin/mount, /bin/umount, /bin/mkdir, /bin/rmdir
```

Examples of Limiting `sudo` Access for the Delphix OS User

In situations where security requirements prohibit giving the Delphix user root privileges to mount, unmount, make directory, and remove directory on the global level, it is possible to configure the `sudoers` file to provide these privileges only on specific mount points or from specific Delphix Engines, as shown in these two examples.

i The Delphix Engines tests its ability to run the `mount` command using `sudo` on the target environment by issuing the `sudo mount` command with no arguments. Many of the examples shown in this topic do not allow that. This causes a warning during environment discovery and monitoring, but otherwise does not cause a problem. If your VDB operations succeed, it is safe to ignore this warning.

However, some users configure the security on the target environments to monitor `sudo` failures and lock out the offending account after some threshold. In those situations, the failure of the `sudo` commands might cause the **delphix_os** account to become locked. One work-around for this situation is to increase the threshold for locking out the user account. Another option is to modify `/etc/sudoers` to permit the **delphix_os** user to run `mkdir`, `rmdir`, `umount` and `mount` command without parameters.

⚠ Note that the following examples are for illustrative purposes and the `sudo` file configuration options are subject to change.

Example 1

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/db2`.

Note that wildcards are allowed for the options on `mount` and `umount` because those commands expect a fixed number of arguments after the options. The option wildcard on the `mount` command also makes it possible to specify the file-system being mounted from the Delphix Engine.

However, wildcards are not acceptable on `mkdir` and `rmdir` because they can have any number of arguments after the options. For those commands, you must specify the exact options (`-p`, `-p -m 755`) used by the Delphix Engine.

Example `/etc/sudoers` File Configuration on the Target Environment for `sudo` Privileges on the VDB Mount Directory Only (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount * /db2/*, \
/bin/umount * /db2/*, \
/bin/umount /db2/*, \
/bin/mkdir -p /db2/*, \
/bin/mkdir -p -m 755 /db2/*, \
/bin/mkdir /db2/*, \
/bin/rmdir /db2/*
```

Example 2

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/db2`, restricts the mount commands to a specific Delphix Engine hostname and IP, and does not allow user-specified options for the `umount` command.

This configuration is more secure, but there is a tradeoff with deployment simplicity. This approach would require a different `sudo` configuration for targets configured for different Delphix Engines.

A Second Example of Configuring the /etc/sudoers File on the Target Environment for Privileges on the VDB Mount Directory Only, and Allows Mounting Only from a Single Server (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount <delphix-server-name>* /db2/*, \
/bin/mount * <delphix-server-name>* /db2/*, \
/bin/mount <delphix-server-ip>* /db2/*, \
/bin/mount * <delphix-server-ip>* /db2/*, \
/bin/mount "", \
/bin/umount /db2/*, \
/bin/umount * /db2/*, \
/bin/mkdir [*] /db2/*, \
/bin/mkdir /db2/*, \
/bin/mkdir -p /db2/*, \
/bin/mkdir -p -m 755 /db2/*, \
/bin/rmdir /db2/*
```

Related Topics

- [Db2 Support and Requirements](#)

Managing Db2 Environments and Hosts


This section covers the following topics:

- [Setting Up Db2 Environments: An Overview](#)
- [Adding a Db2 Environment](#)
- [Managing Db2 Instances](#)
- [Managing Db2 Users and Instance Owners](#)
- [Environment Attributes for Hosts with Db2](#)

Setting Up Db2 Environments: An Overview

This topic describes the high-level process for adding Db2 environments, linking Db2 instances to the Delphix Engine, and provisioning virtual databases.


Delphix for Db2 Architecture



Delphix uses a Standby server model along with Db2s High Availability Disaster Recovery feature to ingest data and stay in sync with the source database. The standby server is then snapshotted by the Delphix Engine and the snapshots can be provisioned out to one or more target servers.

⚠ The snapshot and provision process occurs on the instance level, all databases that exist on the standby server will be provisioned out to the target machines. Similarly, actions such as bookmark, rewind, and refresh will simultaneously apply to all the databases in the instance.

Block Diagram of Linking Architecture Between Db2 Environments and the Delphix Engine



The linking process sets up a database inside an instance on the standby server and uses this as a HADR standby for the primary database. The staging instance must have access to a recent backup copy of the database that will be used to restore the initial copy of the dSource. Once the restoration process is complete, Delphix will issue the


HADR standby commands for the given database and ensure that the health of the HADR connection stays within the acceptable threshold values.

Db2 Staging Instance Set Up

The database level feature enables you to have 1-many mappings between instance and databases. It is also possible to set up databases from multiple primary hosts to use the same standby instance.

The choice of databases on the staging server should also take into account the expected network traffic that HADR will create between the source and staging environments.

Db2 DPF Block Diagram



While using DPF, the DPF environment for Db2 staging and target should be configured on separate hosts. Also, staging and target should be configured with the same number of logical partitions as the number of partitions in the source node.

Related Topics

- [Quick Start Guide for Db2](#)
- [Managing Db2 Environments and Hosts](#)

Adding a Db2 Environment

This topic describes how to add a Db2 Staging environment.

Prerequisites

- Make sure that the staging environment in question meets the requirements described in [Requirements for Db2 Hosts and Databases](#)

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions (...)** menu and select **Add Environment**.
5. In the Add Environment dialog, select **Unix/Linux** in the menu.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter a **Name** for the environment.
9. Enter the **Host IP address** or **hostname**.
10. Enter the **SSH port**.
The default value is 22.
11. Enter an **OS Username** for the environment.
For more information about the environment user requirements, see [Requirements for Db2 Hosts and Databases](#).
12. Select Login Type.
 - Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

 Currently, the Enterprise Password Vault option is not supported.

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For Password, enter the password associated with the user in step 11.

14. Using Public-Key Authentication:
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write privileges to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.
15. You can also add public key authentication to an environment user's profile by using the command-line interface, as explained in the topic [CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users](#).
16. For Password Login, click Verify Credentials to test the username and password.
17. Enter a **Toolkit Path** (make sure the toolkit path does not have spaces).
For more information about the toolkit directory requirements, see [Requirements for Db2 Hosts and Databases](#).
18. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
19. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
20. Click **Submit**.

As the new environment is added, you will see two jobs running in the Delphix platform Job History, one to Create and Discover an environment, and another to Create an environment. When the jobs are complete, you will see the new environment added to the list in the Environments tab. If you do not see it, click the Refresh icon in your browser.

Post-Requisites

To view information about an environment after you have created it:

1. Click **Manage**.
2. Select **Environments**.
3. Select the **environment name**.

Related Topics

- [Setting Up Db2 Environments: An Overview](#)


Managing Db2 Instances

This topic explains how to manage DB2 instances on a host.

When you add an environment with the Delphix Management Application, all existing DB2 instances on the host are automatically discovered by Delphix. A list of all instances and databases are available to Delphix based on the environment discovery process.

View Instances

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, click on the **name of the environment** to you want to refresh.
5. Select the **Databases** tab to see a list of all DB2 instances found in the environment.



The screenshot displays two instance cards in a list. Each card has a header with a dropdown arrow and the instance name and version (e.g., 'auto1054 - 10.5.0.5 -'). Below the header, there are two tabs: 'Installation Details' (selected) and 'Databases'. The 'Installation Details' tab shows various properties for each instance, including 'Allow Provisioning' (Yes), 'Version' (10.5.0.5), 'Install Path' (/opt/ibm/db2/V10.5), 'Instance Path' (/home/auto1054/), 'name' (auto1054), 'gid' (102), 'Pretty Name' (auto1054 - 10.5.0.5 -), 'Instance Name' (auto1054), and 'uid' (533). The 'Databases' tab is active and shows the message 'No databases found on installation.' with a plus sign icon in the top right corner of the tab area.

Updating Instances on a Host

If you add or remove any DB2 instances from a host you must refresh the environment using the steps listed in [Environment Refresh](#) to update the list of instances available to Delphix. Additionally, the environment users must be updated to match the DB2 instance owners as detailed in [Managing DB2 Users and Instance Owners](#).

Related Topics

- [Managing Db2 Users and Instance Owners](#)

Managing Db2 Users and Instance Owners

All DB2 Delphix operations such as linking, snapshot, and provisioning always occur on the database level. As a result, we require that any DB2 instance that you wish to use either as Staging or Target must have its associated instance owner added as an environment user to the host within Delphix. Please perform the following steps to add each DB2 instance owner you to use on the host.

It is important to note that Delphix will not create new instances on the host and can only use existing DB2 instances that it finds on the host machine. For more information on instance management refer to [Managing Db2 Instances](#).

Adding Environment Users

Procedure

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. Click on the existing environment name you want to modify and open the environment information screen.
5. In the Details tab, click the **Plus** icon located next to Environment users.
6. Enter the Username and Password for the OS user in that environment.

- i** If you want to use public-key encryption for logging into your Unix-based environment:
- a. Select **Public Key** for the **Login Type**.
 - b. Click **View Public Key**.
 - c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 authorized_keys` to enable read and write privileges for your user.
 - ii. Run `chmod 755 ~` to make your home directory writable only by your user.
- The public key needs to be added only once per user and per environment.

7. Click the **Check** icon to save the new user.
8. To change the primary user for this environment, click the **Star** icon next to Environment Users. Only the primary user will be used for environment discovery.
9. To delete a user, click the **Trash** icon next to their username.

Related Topics

- [Managing Db2 Instances](#)
- [Managing Db2 Environments and Hosts](#)

Environment Attributes for Hosts with Db2

This topic describes the more advanced attributes for Db2 data platforms.

Common Editable Attributes

Attribute	Description
Environment Users	The users for that environment. These are the users who have permission to ssh into an environment or access the environment through the Delphix Connector. For more information on the environment user requirements, see the Requirements topics for specific data platforms.
Host Address	The IP address of the environment host.
SSH Port	The SSH port number used for the environment.
Toolkit Path	The plugin path for the environment. Delphix uses the Toolkit path (Toolkit is a synonym for Plugin) for logs, mount points, and keeping lib files.
Notes	Any other information you want to add about the environment

Related Topics

- [Managing Db2 Environments and Hosts](#)

Linking Data Sources and Syncing Data with Db2

This section covers the following topics:

- [Linking a dSource from a Db2 Database: An Overview](#)
- [Linking a Db2 dSource](#)
- [Data Management Settings for Db2 dSources](#)
- [Backup Source Database](#)
- [Working with Db2 Snapshots](#)

Linking a dSource from a Db2 Database: An Overview

This topic describes basic concepts behind the creation of dSources from Db2 instances.

Database Level Operation

Associated with an instance is the concept of an instance owner. This is the user who "owns" that instance and has SYSADM authority over the instance and all databases inside that instance. SYSADM authority is the highest level of authority in Db2 and lets this user perform several database management activities such as upgrade, restore, and editing configurations. More information about instances is located in the [IBM knowledge center](#).

Db2 Database Level Support

Delphix operates at the database level and requires that

- The staging and target hosts must have the empty instances created prior to Delphix using them
- The desired OS user to execute commands related to dSource and VDB operation for each instance has been added as an environment user

Note: Make sure the staging instance used for linking doesn't have an existing restoring database with the same name.

Data Ingestion

Db2 for Delphix ingests data by using a staging database created on a discovered Standby instance of Db2. The dSource uses the staging database to stay in sync with the production database. This is done by first going through the linking process during which a full backup is used to recover an initial copy of the production database to the staging database. Storage for the staging database is provided by Delphix via an NFS mount to storage exported by the Delphix system.

 A single standby instance can contain data from multiple source databases.

Pipelining Logic for Implementing Parallel Restores

The plugin employs a parallel pipeline methodology so that the restore operation of non-catalog partitions can be performed in parallel in the Database Partitioning Feature (DPF). The number of parallel restores is determined by the value of "restorePipelineLimit" (default value is 10) in <Toolkit directory>/advanceConfFlag.conf. For more details on this please refer to [Requirements for Db2 Hosts and Databases](#). The parameter "restorePipelineLimit" is configurable by end-users. The plugin performs parallel restore for all non-catalog partitions. For e.g. If the total number of non-catalog partitions is 15, and the "restorePipelineLimit" parameter is set to 10, the first set of 10 restores will happen in parallel. The plugin will track the restore of each partition present in the pipe. Whenever a restore of a partition completes, it will move out from the pipe and a new partition will enter into that pipe. Thus, the plugin ensures that the pipe will always have the user-configured number of partitions being restored (default=10).

In case of a restore failure for a particular partition, the plugin will track that failure and will print the same in the plugin-generated logs so that the user can take the necessary action. If the user received failure during restore operation for a few partitions and then decided to manually initiate the restores for failed partitions, the plugin will verify the datapaths during snapshot operation. In case it finds wrong path information (such as paths outside the mount point), then the plugin will error out of the operation.

The plugin will track the status of each restore in a file named <DB Name>_metadata.json. The plugin can handle a scenario where the staging host reboots in-between the parallel restores. In this case, the user can perform a

“Resynchronisation” operation again. At this time, the remaining restores will resume from scratch. For more details on this please refer to [Requirements for Db2 Hosts and Databases](#).

Data Synchronization

Delphix provides the following options to keep your dSource in sync with the source database:

1. During the linking process, you can set up a HADR connection between the original source database and the Standby instance. This allows the Standby instance using HADR for log shipping to always keep its databases in sync with the source. It is important to note that a single Standby instance (dSource) can contain multiple databases from multiple different servers and instances as long as each database has a unique name.
2. The users can opt for the "**Staging Push**" feature by selecting the checkbox "**Use Staging Push**". With this feature, the user can manage the restore and roll forward of a staging database without involving the Delphix Engine.
3. Users can opt for a non-HADR method of database ingestion. The database will be ingested to a standalone staging host and the user can trigger a resynchronization operation with a most recent online full backup to keep the staging database up to date with the source database.
4. Users can opt for the Delphix Recovery with the **Customer Supplied Archive Logs** method. During an update, the dSource user can trigger a snapshot operation where user-provided archive logs will get applied to the staging database.
5. Optionally, when the DPF feature is in use, users can use the "**Customer Supplied Archive Logs**" feature. Users need to place archive logs inside a folder with a name as NODE<Partition number> where <Partition number> is a four-digit number. For example, if the source environment has two partitions then the user-provided log path will have folder names NODE0000 and NODE0001. Both the folders will have respective archive logs. Snapshot operation will be used to apply the archive logs.



Delphix HADR standby configuration may not meet Db2 DR requirements and should only be used for Delphix use cases.

Related Topics

- [Linking Data Sources and Syncing Data with Db2](#)

Linking a Db2 dSource

This topic describes how to link a Db2 staging dSource.

Prerequisites

- Be sure that the source and staging instances meet the host requirements and the databases meet the container requirements described in [Requirements for Db2 Hosts and Databases](#).

Source Database Preparation

- Instance Owner Permissions

Delphix uses the Db2 instance owner account on the dSource for multiple operations, including verifying the data inside the databases. In order to ingest a database on the staging server with different instances, we need permissions on the source database in order to perform a restore on the staging server. As an example, if we have an instance named "auto1051" at the source with a database name "delphix" and we want to create a dSource on the "auto1052" instance on the staging server, then you must explicitly grant DBADM and SECADM to the dSource instance "auto1052" on the source instance using the following steps:

- Connect to the source database as the source instance owner.
 - connect to <DB_NAME> user <SOURCE_INSTANCE_OWNER>*
- Issue database grant command
 - grant DBADM, SECADM on the database to user <DSOURCE_INSTANCE_OWNER>*
 - grant DBADM, SECADM on the database to user <VDB_INSTANCE_OWNER>*
- Repeat step 2 for every database to be included in the dSource, on the corresponding source database.

Determine if your dSource will be a non-HADR instance, a HADR single standby instance, or a HADR multiple standby instance. Non-HADR dSources can only be updated via a full dSource resync from a newer backup file

Non-HADR Database

- See the "InstanceOwner Permissions" section above.
- Uncheck the default snapsync policy during the dsource linking operation
- Ensure that the source database has the necessary user permissions for the provisioned VDBs as described in [Database Permissions for Provisioned Db2 VDBs](#).

HADR Single Standby Database

- All items in Non-HADR Database section above.
- The following database configuration settings must be set:
 - update db cfg for <DB_NAME> using HADR_LOCAL_HOST <PRIMARY_IP> HADR_LOCAL_SVC <PRIMARY_PORT> immediate*
 - update db cfg for <DB_NAME> using HADR_REMOTE_HOST <STANDBY_IP> HADR_REMOTE_SVC <STANDBY_PORT> immediate*
 - update db cfg for <DB_NAME> using HADR_REMOTE_INST <STANDBY_INSTANCE_NAME> immediate*
 - update db cfg for <DB_NAME> using HADR_SYNCMODE SUPERASYNC immediate*
- If database configuration parameter LOGINDEXBUILD is set to OFF, do the following:
 - update db cfg for <DB_NAME> using LOGINDEXBUILD ON*
 - Force off all connections to the database and reactivate the database
- If database configuration parameter LOGARCHMETH1 is set to OFF, do the following:
 - update db cfg for <DB_NAME> using LOGARCHMETH1 XXXX (must be a valid log archiving method)*

- ii. Take an offline backup
- e. If LOGARCHMETH1 points to a third-party backup server (i.e. TSM or Netbackup) define LOGARCHMETH2 to disk
 - i. *update db cfg for <DB_NAME> using LOGARCHMETH2 DISK:<full path to archive log directory>*
 1. Log files in the directory must be available from the time of the backup until the restore has successfully completed on the dSource.
- f. *db2 start hadr on db <DB_NAME> as primary by force*
- g. Take a full online backup as defined in the "Backup Source Database" section below.
- h. Record the following information, as it must be entered on the Delphix Engine while creating the dSource.
 - i. HADR Primary hostname
 - ii. HADR Primary SVC
 - iii. HADR Standby SVC (auxiliary standby port)

HADR Multiple Standby Databases

This **assumes** a single standby database HADR setup already exists. The existing standby will be referred to as the main standby. The new delphix standby will be referred to as the auxiliary standby.

- a. The following database configuration settings must be set on the primary database:
 - i. *update db cfg for <DB_NAME> using HADR_SYNCMODE <SYNC MODE> immediate* – set whichever sync mode you wish to use on your main standby.
 - ii. *update db cfg for <DB_NAME> using HADR_TARGET_LIST "<MAIN_STANDBY_IP:MAIN_STANDBY_PORT|AUXILIARY_STANDBY_IP:AUXILIARY_STANDBY_PORT>" immediate*
 1. You may have up to two auxiliary standbys defined separated by a '|'; one of which must be the delphix dSource.
- b. *stop hadr on db <DB_NAME>*
- c. *start hadr on db <DB_NAME> as primary by force*
- d. Take a full online backup as defined in the "Backup Source Database" section below. While this backup is running, you may continue with step 5.
- e. The following database configuration settings must be set on the existing main standby database:
 - i. *update db cfg for <DB_NAME> using HADR_SYNCMODE <same mode as defined in 1.a above.>* – It must be the same value used for the primary database.
 - ii. *update db cfg for <DB_NAME> using HADR_TARGET_LIST "<PRIMARY_IP:PRIMARY_PORT|MAIN_STANDBY_IP:MAIN_STANDBY_PORT>"*
- f. *stop hadr on db <DB_NAME>*
- g. *start hadr on db <DB_NAME> as standby*
- h. Record the following information, as it must be entered on the Delphix Engine while creating the dSource (the auxiliary standby database):
 - i. HADR Primary hostname
 - ii. HADR Primary SVC
 - iii. HADR Standby SVC (auxiliary standby port)
 - iv. *HADR_TARGET_LIST <PRIMARY_IP:PRIMARY_PORT|MAIN_STANDBY_IP:MAIN_STANDBY_PORT>*

Delphix Recovery with Customer Supplied Archive Logs

The Delphix Recovery with Customer Supplied Archive Logs (formerly known as Delphix Backup and Logs ingestion), feature allows users to provide Delphix with the location of the source database archive logs. Full backups only have to be restored during dSource creation, or during RESYNC operations when there are exceptions, for example, lost/corrupt log files. Archived logs get applied during the snapshot operation.

A full backup is required to create a dSource and prior to the snapshot operation, users can provide the archive logs to update the dSource during the snapshot operation, user-provided logs will be applied to the staging database.

Applying the logs on a regular basis allows users to keep their dSource in sync with their source database. The staging database will always be in a roll forward pending state. The Plugin performs validation before applying the logs. Plugin managed Pre and Post Hooks can be used for preparing and purging the full backups and archive logs. A list of logs that are available for purging is stored in the snapshot metadata of the dSource. In case DPF is enabled, users need to place archive logs inside a folder with a name as NODE<Partition number> where <Partition number> is a four-digit number. For example, if the source environment has two partitions, then the user-provided log path will have folder names as NODE0000 and NODE0001. Both these folders will have respective archive logs. Snapshot operation will be used to apply the archive logs.

HADR and non-HADR configured dSources can be moved to this feature by performing a RESYNC operation.

Delphix supports Customer Supplied Archive Logs with Db2 version 10.5 and 11.1. To use Customer Supplied Archive Logs simply check the checkbox provided during dSource creation. Customer Supplied Archive Logs require that:

- A full backup must be accessible on the staging server, it could be a shared location but must be mounted on the staging host and visible as a file system.
- Archived logs must also be accessible on the staging server
 - The following command can be used to generate archive logs from the source database;
 - `db2 archive log for database <database name>`
 - Instance users must have read permissions on logs so that Delphix-managed recover can apply the logs.
 - Instance users must have read permissions on logs and additional write permissions on customer-provided log location if Delphix is planning to use customer-provided log location for applying the logs.
 - If the user has placed the archive logs at a user-provided log location, then the logs must be valid and the first active log of the staging database must be present in the user-provided log location.
- If database configuration parameter LOGINDEXBUILD is set to OFF, do the following:
 - `update db cfg for <DB_NAME> using LOGINDEXBUILD ON`
 - Force off all connections to the database and reactivate the database
- The customer does log purging from the provided log location. Delphix only provides the list of logs that are eligible for purging.

Limitations:

- Continuous updates to the dSource are not available. The DB_SYNC operations apply new logs from the log location and keep the dSource updated with changes on the source.
- Read access to the staging database is not available.

Staging Push

With the "Staging Push" feature, the restore and rollforward operations will be managed by the user from the CLI. During dSource creation (or RESYNC) the plugin will create an empty mount point and will provide a restore and rollforward template that will help the user to restore the database. The users can restore the dSource using the steps mentioned in the readme or through any other method of their choice. Once the dSource is ready, the user can perform restore and rollforward operations from the CLI. After the restore, the staging database should remain in the rollforward pending state, and prior to the next snapshot, the user can do a rollforward to a specific Point-in-Time (by providing a timestamp to the rollforward command) against the available archive logs.

Prior to the snapshot operation, the user should validate the backups and the archive logs before applying them to the staging database. Below are the user side validations:

User side validation prior to snapshot operation

- The user must validate the backup files using **db2ckbkp** utility on the staging host.
- The user must validate the archive logs using **db2cklog** utility on the staging host.

- The authenticity of archive logs would be managed by the user.
- The user has to provide the first active log at the staging database log directory location after applying the logs.
- If there is a scenario where a user performed some load copy (non-logged transactions) operations on the source side then it's the responsibility of the user to make those non-logged transactions available on the staging database.

Backup Source Database

New Feature: Source Database with Raw DEVICE type Storage

Several users use raw device-based tablespaces for source Db2 databases. To leverage these environments with Delphix, Delphix has built a workflow using Db2s native tools that allow Delphix to discover and convert a raw device-based tablespace into an automatic storage-based tablespace during ingestion. Once the data is ingested into staging, customers will be able to provision VDBs of the automatic storage-based database.

In order to complete the linking process, the Standby dSource must have access to a full backup of the source Db2 databases on disk. This should be a compressed online Db2 backup and must be accessible to the dSource instance owner on disk. Delphix is currently not set up to accept Db2 backups taken using third-party sources such as Netbackup or TSM. All three features of data ingestion, namely HADR, Non-HADR, and Customer Supplied Archive Logs backups must also include logs. Starting with the Db2 plugin version 3.1.1, we are also supporting full online backups, which are taken using the "exclude logs" syntax. This support is only applicable to the dSources created using the Customer Supplied Archive Logs (Backup and Logs) ingestion mechanism on a single partition.

Support for named pipe type DB2 backup

Users can also provide Db2 backups of type named-pipe. Db2 plugin will check the type (type will be either named-pipe or a file) of backup and will use the restore syntax accordingly.

Example backup command: *db2 backup database <DB_NAME> online compress include logs.*

Backup types supported in DPF

The type of backups supported by Db2 plugin when used with DPF are:

- SSV backups through backup files or named pipes.
- Non-SSV backups through backup files or named pipes.

Example backup command for DB2 SSV backups: `db2_all "db2 BACKUP DATABASE<DB_NAME> ON all dbpartitionnums online to <Backup File Location> compress include logs"`

Example backup command for DB2 Non SSV backups: `db2_all "db2 backup database <DB_NAME> online to <Backup File Location> compress include logs"`

Best Practices for Taking a Backup

The following best practices can help improve backup and restore performance:

1. Compression should be enabled
2. Following parameters should be optimally configured:
 - a. Utility Heap Size (UTIL_HEAP_SZ)
 - b. No. of CPUs

- c. No. of Table Spaces
- d. Extent Size
- e. Page Size

1. Parallelism & Buffer configuration may be used to improve backup performance. Parameters that should be configured are:
 - a. Parallelism
 - b. Buffer Size
 - c. No. of Buffers

More information about backup best practices is available in IBM Knowledge Center.

Procedure

1. Login to the **Delphix Management Application**.
2. On the **Databases** tab of the Environment Management screen, add a source config against the discovered staging instance.
3. Then, click **Manage** and select **Datasets**.
4. Click the **Plus (+)** icon and select **Add dSource**, you'll get a list of available source configs using which you can go for dSource creation.
5. In the **Add dSource** wizard, select the required source configuration. Consider the following when choosing an ingestion method.
 - a. Users can specify their own mount path to host the dSource dataset by specifying the path in the **Mount Base** field.
 - b. The restore and rollforward operations on the staging database can be managed by selecting the **"Use Staging Push"** checkbox. If the **"Use Staging Push"** checkbox is selected, the other checkboxes must not be selected.
 - c. If you are working with a HADR setup and want to ingest the database on a staging using HADR configurations, you need to check the Link against source using HADR checkbox and need to uncheck the Customer Supplied Archive Logs checkboxes.
 - d. To use a non-HADR method for database ingestion user need to uncheck the **Link against the source using HADR** and **Customer Supplied Archive Logs** checkboxes.
 - e. To use the **Customer Supplied Archive Logs** method for database ingestion check the **Customer Supplied Archive Logs** checkbox and uncheck the **Link against source using HADR** checkbox and enter the path in **Archive Log Path**.
 - f. To use the **Database partition feature** check **Use DPF** checkbox and uncheck the **Link against the source using HADR** checkbox. For more information on DPF please view [IBM Db2 Overview](#).

Source

Select an available environment /data source to connect to your dSource

Data Source

Filter: none ▼

☰ **R82D115A**
 new82.dcol1.delphix.com
 AppData

Data Type

AppData

Database Name (Required) *

Database name to restore from.

Mount Base

Path on the local filesystem where the dSource will be mounted by the Delphix Engine. The new path will take into effect only after executing the disable and enable operations.

Staging Push

Enabling this option will allow user to manage the staging database themselves. If this option is selected then there is no need to provide backup path and no need to select any other option.

Backup Path

Location of backup files on the Staging host. Staging host selection will be required before completing this wizard. Default backup file location is the instance home directory.

DPF (Database partitioning feature)

Enabling this option will allow creation of a dSource with DPF (Database partitioning feature) configuration

Archive Logs

Link a database, optionally providing the location of Archive Logs.

Archive Log Path

Location of Customer Supplied Archive Logs which will be applied to the staging database prior to any snapshot. Path must be present on the Staging host.

Archive Log Location is READ ONLY

Allow Delphix to copy archive logs to a temporary location, apply the logs, remove the temporary log store and then snapshot the updated staging database. This strategy may be applicable where archive logs are shared from production over NFS and the preference is to set the share to RO.

Rollforward to PIT

Rollforward a database to a particular timestamp if archive logs are provided and a valid timestamp is provided along with snapshot parameters.

Link against source using HADR

Enabling this option will configure the staging host as a HADR Standby.

HADR Primary Hostname

The host name where the primary database resides (required for HADR databases).

HADR Primary SVC

HADR Primary SVC port value (required for HADR databases).

HADR Standby SVC

HADR Standby SVC port value (required for HADR databases).

HADR Target List

HADR Target list 'host:port|host:port|host:port' values (required for HADR with multiple standby databases).

Use Plugin Hooks

Execute custom scripts when linking a dSource

Plugin Pre-Restore Hook

Execute custom scripts prior to any recovery operation.

Plugin Post-Restore Hook

Execute custom scripts following database recovery as a part of linking or resync operations.

Staging Database Configuration Settings

Customise Staging Database configuration settings. For example DATABASE_MEMORY, LOGFILSIZ.

+ Add

Cancel
Back
Next
Submit

6. The **database name** is mandatory and must be unique for a given instance. This is the name that the database was on the instance it was restored from.
7. If the user is planning to create a dSource with the "**Staging Push**" feature then the user must need to select the "**Use Staging Push**" checkbox.
8. Enter the complete **Backup Path** where the database backup file resides. If no value is entered, the default value used is the instance home directory. If there are multiple backup files for a non-DPF database on the backup path, the most current one will be used. If there are multiple sets of backup files for a DPF database on the backup path, the most recent one will be used.
9. If the user wants to roll forward a snapshot to a user-specified timestamp, then the user must select the **Rollforward to PIT** checkbox. Users can also manually do this through the snapshot option.

 The user-specified timestamp needs to be provided along with the timezone. For example, 2021-11-14-04.15.00 **UTC**. If the timezone is not entered, then the user's local timezone is considered.

 Uncheck the default snapsync policy during the dsource linking operation.

Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section.

10. Optionally, users can provide a bash script as a **Plugin** defined hook. Delphix has two Plugin managed Hooks:
 - a. **Plugin Pre-Restore Hook:** The bash script entered in this field will get executed before restoration. This hook is useful in such cases where the users want to prepare the full backup before the restore operation.
 - b. **Plugin Post-Restore Hook:** The bash script entered in this field will get executed after completion of the restore operation. This hook is useful in such cases where the user wants to purge the full backup file after the restore operation.

1. Plugin managed pre-Hook - User can copy the required full backup file from some remote location to staging host. Below is an example of such bash script (name of below script is copyBackup.sh)

```
#!/bin/bash
```

```
# Copying backup file from a remote host remote.delphix.com
scp auto1051@remote.delphix.com:/backuplocation/
R74D105A.0.auto1051.DBPART000.20180510073639.001 /db2backup/
```

The user can provide the above script in Plugin managed pre-Hook field as "sh copyBackup.sh"

2. Plugin managed post-Hook - User can purge the full backup file from the staging host. Below is an example of such bash script (name of below script is purgeBackup.sh)

```
#!/bin/bash
```

```
# Purging full backup file from staging host for saving the staging storage
rm -f /db2backup/R74D105A.0.auto1051.DBPART000.20180510073639.001
```

The user can provide the above script in Plugin managed post-Hook field as "sh purgeBackup.sh"

11. If the dSource is to use **HADR** please enter the following fields. For more information about HADR please view [Linking a dSource from a Db2 Database: An Overview](#).
 - a. Enter a fully qualified HADR Primary Hostname. This is a required field for HADR and must match the value set for HADR_LOCAL_HOST on the master.
 - b. Enter the port or /etc/services name for the HADR Primary SVC. This is a required field for HADR and uses the value set for HADR_LOCAL_SVC on the master.
 - c. Enter the port or /etc/services name for the HADR Standby SVC. This is a required field for HADR and uses the value set for HADR_REMOTE_SVC on the master.
12. If you are using **Customer Supplied Archive Logs** for dSource creation complete the following fields:
 - a. Check the Customer Supplied Archive logs checkbox
 - b. Provide the log location of the archive logs. In this location, users can load copy files by appending schema and table name to it so that plugin can process those load copy files during log apply process. Archive logs will be used for applying the logs during the snapshot operation. This location must be a part of the staging server and this location must be different than full backup location.

The screenshot shows the 'Add dSource' configuration window. On the left, a sidebar lists steps: Source (selected), dSource Configuration, Data Management, Policies, Hooks, and Summary. The main content area is divided into two sections. The top section is for 'Customer Supplied Archive Logs', with a checkbox that is unchecked and a text input field for 'Archive Log Path'. The bottom section is for 'Archive Log Location is READ ONLY', with a checkbox that is checked. Below this are several text input fields for HADR-related parameters: 'HADR Primary Hostname', 'HADR Primary SVC', 'HADR Primary SVC port value', 'HADR Standby SVC', 'HADR Standby SVC port value', and 'HADR Target List'.

- c. The **Archive Log Location is READ ONLY** (default) option indicates whether Delphix should copy the log files to the dSource’s filesystem prior to the log apply process. If this is not selected the staging database will read the log files directly from the user-provided log location and perform the log apply operation.

⚠ Special considerations for Archive Log Location is READ ONLY

If this operation is performed on the source database and user want the Plugin to process the load files place the load copy files in your location and append the schema and table names (for example: <DB Name>.<Number>.<Instance>.DBPART<partNumber>.<Timestamp>.001.<schema name>.<table name>) to the load copy file names
If you do not want to process load copy files via the Plugin place the load copy files in a similar structure as it is on the source or set a DB2 registry parameter DB2LOADREC.

⚠ Special consideration during dSource upgrade from older version plugin to Customer Supplied Archive Logs feature enabled plugin

Suppose the user is having a dSource with an older version of the plugin and now wants to upgrade that dSource to **Customer Supplied Archive Logs** feature then the user needs to perform the following steps :

1. Go to the configuration tab of dSource and then click on the custom sub-tab.
2. Click on the edit icon.
3. Check the **Customer Supplied Archive Logs** and **Archive Log Location is READ ONLY** checkbox.
4. Click on submit and then go for RESYNC operation to upgrade the dSource.

13. Click **Next**.
14. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
15. Click **Next** and set the **Staging Environment** to be the same as the dSource host. Select the **Staging Environment User** to be the same as the instance owner of the dSource instance.

✓ **Changing the Environment User**

If you need to change or add an environment user for the dSource instance, see [Managing Db2 Users and Instance](#).

⚠ Ensure that the Staging environment must be the environment that you have chosen to create the source config.

Select the **User** from the drop-down and set the **SNAPSHOT PARAMETERS** for the dSource snapshot.

- a. Click on the **Timestamp** input field to enter the timestamp at which the snapshot will be taken. The user-specified timestamp needs to be provided along with the timezone in YYYY-MM-DD-HH.MM.SS format. For example, 2022-03-12-04.15.00**UTC**. If the timezone is not entered, then the user's local timezone is considered.

⚠ Providing the timestamp will only work if the user has checked the **Rollforward to PIT** option in step 9 above.

- b. Select **Resynchronize dSource** to resynchronize the dSource. This will force a non-incremental load of data from the source similar to when the dSource was created. This action avoids duplicating storage requirements and maintains a Timeflow history.
16. The Delphix Engine will initiate two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking Active Jobs in the top menu bar, or by selecting System > Event Viewer. When the jobs have completed successfully, the database icon will change to a dSource icon on the Environments > Host > Databases screen, and the dSource will also appear in the list of Datasets under its assigned group. Then, click **Next** and you'll get the Policies section. Set the desired **Snapsync Policy** for the dSource. For more information on policies see Advanced Data Management Settings for Db2 dSources.
 17. Click **Next** and specify any desired pre- and post-scripts. For details on pre- and post-scripts, refer to Customizing Db2 Management with Hook Operations.
 18. Click **Next**. Review the dSource Configuration and Data Management information in the Summary section.
 19. Click **Submit**.

ℹ **The dSource Configuration Screen**

After you have created a dSource, the **dSource Configuration tab** allows you to view information about it and make modifications to its policies and permissions. In the **Datasets** panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the **Source files, Data Management** configuration, and **Hook Operations**.

⚠ **Purging of archive logs after snapshot operation**

Logs can be purged from your log location after the snapshot operation. The Plugin will preserve the list of logs that are eligible for purging in the snapshot's metadata. The snapshot metadata can be accessed via the Delphix CLI or API.

Related Topics

- [Sudo Privilege Requirements for Db2 Environments](#)
- [Sudo File Configuration Examples for Db2 Environments](#)

- [Managing Db2 Users and Instance Owners](#)

Data Management Settings for Db2 dSources

This topic describes advanced data management settings for dSources.

When linking a dSource, you can use custom data management settings to improve overall performance and match your specific server and data environment's needs. If no specific settings are required, leverage default data management settings.

Accessing Data Management Settings

There are three ways to set or modify data management settings for dSources:

1. During the dSource linking process, in the Policies tab of the Add dSource wizard, select the **Retention Policy**.
2. Or select **Manage**, then select **Policies**.
3. In the Policies screen, select the **Retention** tab. To manage retention on target after the dSource snapshots have been deleted from the replication source, select the **Replica Retention** tab.
4. To add a new retention policy click the **+Retention** button, and to add a new replica retention policy on the replication target, click the **+Replica Retention** button.
5. This will open the Retention or the Replica Retention Policy wizard, depending upon the above selection. Enter a policy name along with required retention details and click **Submit**.
For more information, see [Policies for Scheduled Jobs](#).

Retention Policies - Retention/Replica Retention

Retention policy defines the length of time Delphix Engine retains snapshots within its storage. To support longer retention times, you may need to allocate more storage to the Delphix Engine. The retention policy – in combination with the SnapSync policy – can significantly impact the performance and storage consumption of the Delphix Engine.

Replica Retention policies define how long the snapshots are retained on replicated namespaces for dSources and VDBs after they have been deleted on the replication source. Normally, the snapshots that have been deleted on the replication source engine are also deleted on the replication target engine. A new retention policy is introduced to provide an extended lifetime of such snapshots on the replication target.

Benefits of Longer Retention

With increased retention time for snapshots and logs, you allow a longer (older) rollback period for your data.

Common use cases for longer retention include:

- SOX compliance
- Frequent application changes and development
- Caution and controlled progression of data
- Reduction of project risk
- Restore to older snapshots

Related Topics

- [Linking Data Sources and Syncing Data with Db2](#)
- [Policies for Scheduled Jobs](#)

Backup Source Database

Source Database with Raw DEVICE type Storage

Several users use raw device-based tablespaces for source DB2 databases. To leverage these environments with Delphix, Delphix has built a workflow using Db2s native tools that allow Delphix to discover and convert a raw device-based tablespace into an automatic storage-based tablespace during ingestion. Once the data is ingested into staging, customers will be able to provision VDBs of the automatic storage-based database.

In order to complete the linking process, the Standby dSource must have access to a full backup of the source Db2 databases on disk. This should be a compressed online DB2 backup and must be accessible to the dSource instance owner on disk. Delphix is currently not setup to accept DB2 backups taken using third-party sources such as Netbackup or TSM. All three features of data ingestion, namely HADR, Non-HADR, and Customer Supplied Archive Logs backups must also include logs. Starting with Db2 plugin version 3.1.1, we are also supporting full online backups, which are taken using the "exclude logs" syntax. This support is only applicable to the dSources created using the Customer Supplied Archive Logs (Backup and Logs) ingestion mechanism on a single partition.

Example backup command: `db2 backup database <DB_NAME> online compress include logs`

Best Practices for Taking a Backup

The following best practices can help improve backup and restore performance:

1. Compression should be enabled
2. Following parameters should be optimally configured:
 - a. Utility Heap Size (UTIL_HEAP_SZ)
 - b. No. of CPUs
 - c. No. of Table Spaces
 - d. Extent Size
 - e. Page Size
1. Parallelism & Buffer configuration may be used to improve backup performance. Parameters that should be configured are :
 - a. Parallelism
 - b. Buffer Size
 - c. No. of Buffers

More information about backup best practices is available in [IBM Knowledge Center](#).

Procedure

1. Login to the Delphix Management Application using Delphix Engine credentials or as the owner of the database from which you want to provision the dSource.
2. On the **Databases** tab of Environment Management screen, add a source config against discovered staging instance.
3. Then, click **Manage**.
4. Select **Datasets**.
5. Click the Plus (+) icon and select **Add dSource**, you'll get a list of available source configs using which you can go for dsource creation.
6. In the **Add dSource** wizard, select the required source configuration.
7. If you are working with an HADR setup, please leave the HADR checkbox checked.

8. The **database name** is mandatory and must be unique for a given instance. This is the name that the database was on the instance it was restored from.
9. Enter the complete **Backup Path** where the database backup file resides. If no value is entered, the default value used is the instance home directory. If there are multiple backup files for a database on the backup path, the most current one will be used.
10. Enter the **Log Archive Method1** you wish to use for the database. If no value is entered, the default value used is `DISK:/mountpoint/dbname/arch`.
11. Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section.
12. If the dSource is to use HADR please enter the following fields. If it will not use HADR skip ahead to step 13. For more information about HADR please view [Linking a dSource from a Db2 Database: An Overview](#).
 - a. Enter a fully qualified HADR Primary Hostname. This is a required field for HADR and must match the value set for HADR_LOCAL_HOST on the master.
 - b. Enter the port or /etc/services name for the HADR Primary SVC. This is a required field for HADR and uses the value set for HADR_LOCAL_SVC on the master.
 - c. Enter the port or /etc/services name for the HADR Standby SVC. This is a required field for HADR and uses the value set for HADR_REMOTE_SVC on the master.
13. Click **Next**.
14. Select a **dSource Name** and **Database Group** for the dSource.
15. Click **Next**.
You will get Data Management section where you need to specify staging environment and user which will be used for dsource creation.
16. Set the **Staging Environment** to be the same as the dSource host.
17. Select the **Staging Environment User** to be the same as the instance owner of the dSource instance.

 **Changing the Environment User**

If you need to change or add an environment user for the dSource instance, see [Managing Db2 Users and Instance Owners](#).

18. Then, click **Next** and you'll get the Policies section. Set the desired **Snapsync Policy** for the dSource.
19. Click **Next**.
20. Specify any desired pre- and post-scripts. For details on pre- and post-scripts.
21. Click **Next**.
22. Review the dSource Configuration and Data Management information in the summary section.
23. Click **Submit**.

The Delphix Engine will initiate two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking Active Jobs in the top menu bar, or by selecting System > Event Viewer. When the jobs have completed successfully, the database icon will change to a dSource icon on the Environments > Host > Databases screen, and the dSource will also appear in the list of Datasets under its assigned group.

 **The dSource Configuration Screen**

After you have created a dSource, the **dSource Configuration tab** allows you to view information about it and make modifications to its policies and permissions. In the **Datasets** panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the **Source files**, **Data Management** configuration, and **Hook Operations**.

Related Topics

- [Linking Data Sources and Syncing Data with Db2](#)
- [Linking a dSource from a Db2 Database: An Overview](#)
- [Managing Db2 Users and Instance Owners](#)

Working with Db2 Snapshots

Taking a snapshot creates a new snapshot entry in the dSource's Timeflow. The plugin-managed dSource allows you to select the parameters before taking the snapshot. These parameters are provided as an input to pre-snapshot and post snapshot functions. On the other hand, in the default snapshot, the parameters are pre-defined in the dSource environment.

Snapshot (Default)

This section lists the steps to take a snapshot and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog, select **Perform Snapshot**.
6. Click **View:** under **Timeflow** to verify the Snapshot you just created.
7. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
8. In the Delete Snapshot dialog, select **Delete**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Provisioning and Managing VDBs from Db2 dSources

This section covers the following topics:

- [Provisioning Db2 VDBs: An Overview](#)
- [Provisioning a Db2 VDB](#)
- [Database Permissions for Provisioned Db2 VDBs](#)
- [Upgrading Db2 VDBs](#)
- [Provisioning from a Replicated Db2 VDB](#)

Provisioning Db2 VDBs: An Overview

This topic describes the basic concepts involved with provisioning VDBs from DB2.

A dSource is a virtualized representation of a physical or logical source database. As a virtual representation, it cannot be accessed or manipulated using database tools. Instead, you must create a virtual database (VDB) from a dSource snapshot. A VDB is an independent, writable copy of a dSource snapshot. You can also create VDBs from other VDBs. Once you have provisioned a VDB to a target environment, you can also implement snapshot and retention policies for the VDB, which will determine how frequently Delphix Engine will take a database snapshot and how long the snapshots will be retained for recovery and provisioning purposes. For an overview of the high-level components involved in provisioning a DB2 VDB refer to [Setting Up Db2 Environments: An Overview](#).

Snapshots accumulate over time. To view a snapshot:

1. From the Datasets panel, click the group containing the dSource.
2. Select dSource.
3. Click the TimeFlow tab.

The TimeFlow appears as a list of dates, each of which expands to show snapshots from that date. Times, when the VDB has been refreshed, are marked by a blue line between dates.

On the TimeFlow, you can also filter by type of snapshot. To do so, click the filter button, which is shaped like an eye.

You can scroll through these lists to select the one you want, or you can enter a date and time to search for a specific snapshot.

Once you have provisioned a VDB, you can also take snapshots of it. As with the dSource snapshots, you can find these when you select the VDB in the Datasets panel. You can then provision additional VDBs from these VDB snapshots.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot.

Related Topics

- [Setting Up Db2 Environments: An Overview](#)
- [Provisioning and Managing VDBs from Db2 dSources](#)

Provisioning a Db2 VDB

This topic describes how to provision a virtual database (VDB) from a DB2 dSource.

Prerequisites

- You will need to have linked a dSource from a staging instance, as described in [Linking a Db2 dSource](#), or have created a VDB from which you want to provision another VDB
- You should have set up the Db2 target environment with necessary requirements as described in [Requirements for Db2 Hosts and Databases](#)
- Make sure you have the required Instance Owner permissions on the target instance and environment as described in [Managing Db2 Users and Instance Owners](#)
- The method for [Database Permissions for Provisioned Db2 VDBs](#) is decided before the provisioning
- [Optional] Once you have restored the databases into your target environment, If the version of IBM® Db2® in your source environment is at a different fix pack level than the version of Db2 in your target environment, then you need to update the 11.1 version databases with db2updv111 command and 11.5 databases with db2updv115 command utility.

Run the below command as a Pre-snapshot hook:

```
db2updv111 -d <DBNAME> -a
db2updv115 -d <DBNAME> -a
```

The optional pre-requisite is applicable only when the target environment is using a different Db2 fix pack level than the source. If your source environment is using version 11.1.x and your target is using version 11.5.x, you can neglect the above-mentioned point.

Consider the following scenarios for the optional prerequisite mentioned above:

- if your source environment uses Db2 version 11.5.5 and your target environment uses Db2 11.5.8 or a later fix pack, you need to update the databases in the target environment.
- If your databases were created or upgraded to Db2 Version 11.1 GA and you applied Db2 Version 11.1 Fix Pack 1 or later, then running the above command will automatically apply all updates required from Version 11.1 GA up to and including the fix pack level that you are installing.



- To know how to update databases, see [db2updv115 - Update database to Version 11.5 mod pack command](#).
- Certain database or instance operations are restricted while an online fix pack update is in progress.
- See [db2updv111 - Update database to Version 11.1 fix pack command](#), which mentions it is mandatory to run db2updv111 on database if one is coming from higher version of fix pack to Lower OR vice versa.
- For more information, you can also refer: [Update the databases for a Db2 11.1.x fix pack](#)



You can take a new snapshot of the dSource by clicking the **Camera** icon on the dSource card. Once the snapshot is complete you can provision a new VDB from it.



Please note while creating a VDB on an instance

By default, DB2 Plugin will not allow provisioning a VDB on an instance which contains a database that has an identical name to the source database.

For e.g.;

1. When we are provisioning a VDB from a VDB then source VDB will be treated as the source database.
2. When we are provisioning a VDB from a dSource then a staging database will be treated as the source database.

If the user wants to provision a VDB on an instance which contains a source database, then the user needs to follow the below steps :

1. Check if any config file with name db2_plugin.conf at location **<Toolkit directory>/Delphix_COMMON/plugin/<Identifier>/**. If it does not exist, then the user should create a file with name db2_plugin.conf at location **<Toolkit directory>/Delphix_COMMON/plugin/<Identifier>/**. The instance user must have read permission to this file. Once the file is created, user should proceed to step 2. If the file already exists at this location with the required permission, the user should directly proceed to step 2.
 2. Add the below line in this config file.
 1. allowSourceDbOnSameInstance=**True**
- Then create/Refresh the VDB again.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource**.
5. Select a **snapshot** from which you want to provision.
6. Click **Provision VDB** icon to open Provision VDB wizard.
7. Select a target environment from the left pane.
8. Select an **Installation** to use from the dropdown list of available DB2 instances on that environment.
9. Set the **Environment User** to be the Instance Owner. Note: The picking of instance owner is only possible if you have multiple environment users set on that host.
10. Provide VDB Name as database name as parameter.
11. Optionally, set the database configuration parameters for the VDB.
12. Click **Next**.
13. Select a **Target Group** for the VDB.
Click the green **Plus** icon to add a new group, if necessary.
14. Select a **Snapshot Policy** for the VDB.
15. Click **Next**.
16. Specify any desired hook operations.
17. Click **Next**.
18. Review the **Provisioning Configuration** and **Data Management** information.
19. Click **Submit**.

When provisioning starts, you can review the progress of the job in the **Databases** panel, or in the **Job History** panel of the **Dashboard**. When provisioning is complete, the VDB will be included in the group you designated and listed in the **Databases** panel. If you select the VDB in the Databases panel and click the **Open** icon, you can view its card, which contains information about the database and its Data Management settings.

Once the VDB provisioning has successfully completed, if the source and target instance ids are not the same, you may want to grant secadm and dbadm on the database to the target instance id. Refer to [Database Permissions for Provisioned Db2 VDBs](#) for more information.

Related Topics

- [Provisioning and Managing VDBs from Db2 dSources](#)

Database Permissions for Provisioned Db2 VDBs

This topic describes the database permissions on provisioned Db2 virtual instances

DB2 Authentication

Authentication is the process of validating a supplied user ID and password using a security mechanism. User and group authentication is managed in a facility external to Db2 LUW, such as the operating system, a domain controller, or a Kerberos security system. This is different from other database management systems (DBMSs), such as Oracle and SQL Server, where user accounts may be defined and authenticated in the database itself, as well as in an external facility such as the operating system.

Any time a user ID and password is explicitly provided to Db2 LUW as part of an instance attachment or database connection request, Db2 attempts to authenticate that user ID and password using this external security facility. If no user ID or password is provided with the request, Db2 implicitly uses the user ID and password that were used to login to the workstation where the request originated. More information on Db2 authentication and authorization is available via [IBM documentation](#).

Delphix Db2 Authentication

Delphix for Db2 requires that the staging and target hosts must already have the necessary users and authentication systems created/installed on them. Delphix will neither create users nor change database passwords as part of the provisioning process.

While the terminology used within the Delphix Management application refers to a VDB, the ingestion, snapshot and provisioning process for Db2 on Delphix always occurs on the Database level. Thus when a virtual Db2 database is provisioned by Delphix, it contains the Db2 database that were in the source instance with the identical user permissions as they had on the source. This means that if the target instance name is different from the source instance name then that instance owner will NOT have DBADM or SECADM permissions unless they were specifically granted to that instance owner on the source instance. The instance owner will however always have SYSADM permissions on all databases in the instance.

LDAP Authentication

If your Db2 instances and applications use LDAP authentication then they will work seamlessly as long as LDAP had been configured on the VDB Target Instance.

OS Authentication

If your DB2 instances and applications are using OS authentication then it is important to ensure that the relevant OS accounts exist on the target machine.

Generic Accounts

If the Db2 applications are using generic (non-instance owner) accounts, they will then be able to continue using them as long as those OS accounts exist on the host machine. It is important to note that the passwords for the same account may be different on different hosts, or if they use different LDAP servers (i.e. prod vs dev LDAP servers).

Instance Owner Usage

If the Db2 applications typically use the instance owner accounts to access data then it is important to note that if the new instance name is different from the source instance name then that instance owner will NOT have DBADM or SECADM permissions in the VDB instance and may not be able to read and write data. In such a case there are three possible workarounds:

1. Connect to the databases as the source instance owner. This requires that the target host has a UNIX account with the same name as the source instance.

2. Grant the necessary permissions to the VDB instance owner on the source instance before taking the snapshot that is to be provisioned.
3. Grant the permissions to a known "delphix" user on the source instance and then use that account to grant the permissions to the target VDB instance after provisioning.

Related Topics

- [Provisioning and Managing VDBs from Db2 dSources](#)

Upgrading Db2 VDBs

This topic describes how to upgrade a Db2 VDB to a higher version of Db2.

 Before upgrading the VDB, please refer to the Delphix Db2 support matrix to ensure that you upgrade to the supported Db2 versions only.

Procedure for VDB In-Place Upgrade

1. Remove any VDB Refresh Policy assigned to the VDB.
2. Upgrade the target Db2 instance.
3. Refresh the target environment.

Related Topics

- [Provisioning and Managing VDBs from Db2 dSources](#)

Provisioning from a Replicated Db2 VDB

This topic describes how to provision from a replicated dSource or virtual database (VDB).

The process for provisioning from replicated objects is the same as the typical VDB provisioning process, except that first you need to select the replica namespace containing the replicated object.

Prerequisites

- You must have a dSource or VDB that has been replicated from one Delphix Engine to another, as described in [Replication Overview](#)
- The Delphix Engine containing the replicated dSource or VDB must have a compatible target environment that it can use to provision a VDB from the replicated dSource or VDB.

Procedure

1. On the Delphix Engine containing the replicated dSource or VDB, login to the **Delphix Management** application.
2. In the top menu bar, click **Manage**.
3. Select **Datasets**.
4. From the list of replica namespaces, select the **replica namespace** that contains the dSource or VDB from which you want to provision.
5. The provisioning process is now identical to the process for provisioning standard objects.

Post-Requisites

Once the provisioning job has started, the user interface will automatically display the new VDB in the live system.

Related Links

- [Db2 on Delphix: An Overview](#)
- [Replication Overview](#)

Db2 Hook Operations

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at `/bin/sh`. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh "$ARG_ENVIRONMENT_VARIABLE"
"second argument in double quotes" 'third argument in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a `bash` binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
```

```
# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

Shell Operation Tips

Using `nohup`

You can use the `nohup` command and process backgrounding from resource in order to "detach" a process from the Delphix Engine. However, if you use `nohup` and process backgrounding, you MUST redirect `stdout` and `stderr`.

Unless you explicitly tell the shell to redirect `stdout` and `stderr` in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either `stdout` or `stderr`. Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your `RunCommand` operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The `RunExpect` operation executes an `Expect` script on a Unix environment. The `Expect` utility provides a scripting language that makes it easy to automate interactions with programs which normally can only be used interactively, such as `ssh`. The Delphix Engine includes a platform-independent implementation of a subset of the full `Expect` functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Management application and CLI to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an `ssh` session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

DB2 Environment Variables

Operations that run user-provided scripts have access to environment variables. For operations associated with specific dSources or virtual databases (VDBs), the Delphix Engine will always set environment variables so that the user-provided operations can use them to access the dSource or VDB.

dSource Environment Variables

Environment Variable	Description
<code>DLPX_DATA_DIRECTORY</code>	The path where the staging database is mounted

VDB Environment Variables

Environment Variable	Description
<code>DLPX_DATA_DIRECTORY</code>	The path where virtual database is mounted

Related Topics

- [Db2 on Delphix: An Overview](#)
- [Db2 Support and Requirements](#)
- [Hook Scripts for Automation and Customization](#)

Enhanced Logging for Db2 Plugin

DB2 Plugin v3.x

This feature adds three options to provide enhanced logging:

1. To set debug log levels
2. To set log retention levels
3. To set trace size limit

Each host (added as an environment to DE) will have its own property file **<Toolkit directory>/DB2/plugin-logging.properties**. Delphix users can modify the parameters within this property file to configure the logging mechanism.

Setting Debug Log Levels

Delphix users can set log levels for plugin generated logs by setting the parameter **level**. Higher logging levels will help to expedite debugging issues. There are three levels of logging which are: Level 1 (info), Level 2 (debug) and Level 3 (trace). All the objects (both dSources and VDBs) present on a staging/target host will refer to a single property file. For example; if we have 4 dSources on a host associated with a Delphix Engine then the log level applies to all the dSources.

Log level description:

- level=1: This level will print only informational logs.
- level=2: This level will print informational logs and debug statements. This is the default log level.
- level=3: Besides informational and debug logs, this level will enable the traces.

Setting Log Retention Levels

Delphix users can set a retention level for diag logs (<Db name>.diag.log) using parameter **retention** in the file. As per this parameter, the log files are moved (archived), renamed or deleted once they are too old or too big. New incoming log data is directed into a new fresh file (at the same location).

By default, this value will be set to a minimum value of 2. The user can change this value and set its value within the range 2 and 50. For example; if **retention** is set to 4, the plugin will have the following log files: <DB Name>.diag.log, <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB Name>.diag.log.3, <DB Name>.diag.log.4.

- File <DB Name>.diag.log is the active log file.
- File <DB Name>.diag.log.1 is the most recent archive log file
- File <DB Name>.diag.log.4 is the oldest one

Setting Trace Size Limit

In case trace level logging is enabled, users can set the size of the trace folder using parameter **traceFolderLimit**. By default, the value of this parameter will be 10, which signifies that the size limit for this folder will be 10 MB. This parameter is configurable and the user can set its value to any positive integer. The traces are stored within a separate folder for traces created under directory **<Toolkit directory>/DB2/logs/<User>/**. The filename format for traces would be <DB Name>.<script name>.trace.<epoch time>.

If the file is accidentally removed from the location, the plugin will generate the property file **<Toolkit directory>/DB2/plugin-logging.properties** with default values at runtime. For example, when the user will upload the plugin to DE, the first call to log function will check for the existence of **<Toolkit directory>/DB2/plugin-logging.properties**. If the plugin does not find this property file, it will create a property file with the content shown below. The logs will

start generating per the level defined inside the file. Below is the sample **<Toolkit directory>/DB2/plugin-logging.properties** file:

plugin-logging.properties

```
# This flag will set the debug level of logging. There are three levels:

# 1: Logging will be in info mode.

# 2: Logging will be in Debug mode.

# 3: Traces will be enabled along with debug level logging.

# The default value is 2.

# level=<positive integer>

level=2

# Whenever the size of <DB Name>.diag.log file exceeds 1MB the plugin will rename the
active log file to <DB Name>.diag.log.<number> and a new log file with name <DB
Name>.diag.log will be generated.

# For example, if LogRetention is set to 4, the plugin will have the following log
files: <DB Name>.diag.log, <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB
Name>.diag.log.3, <DB Name>.diag.log.4.

# File <DB Name>.diag.log.4 will be the oldest one.

# File <DB Name>.diag.log.1 will be the most recent archive log file.

# File <DB Name>.diag.log will be the active log file.

# For example, if LogRetention is set to 4, the plugin will have the following log
files: <DB Name>.diag.log, <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB
Name>.diag.log.3, <DB Name>.diag.log.4.

# The default value for retention flag is 2, minimum value is 2 and maximum value is
50.

# retention=<positive integer>

retention=2

# The value for this property is configurable and must be a positive integer.

# The default value is 10.

# traceFolderLimit=<positive number>

traceFolderLimit=10
```

DB2 plugin v4.x and onwards

This feature adds three options to provide enhanced logging:

1. To set debug log levels.
2. To set active log file size.
3. To set log retention levels.

Delphix users can modify the parameters defined under `plugin_logging_parameters` section header within the `db2_plugin.conf` file located at `<Toolkit directory>/<Delphix_COMMON>/plugin/DB2_18f4ff11-b758-4bf2-9a37-719a22f5a4b8/db2_plugin.conf` to configure the logging mechanism.

Setting Log Levels (`level`)

Delphix users can set log levels for plugin generated logs by setting the parameter ***level*** under `plugin_logging_parameters` header. Higher logging levels will help to expedite debugging issues. There are two levels of logging which are: Info and Debug. All the objects (both dSources and VDBs) present on a staging/target host will refer to a single property file. For example; if we have 4 dSources on a host associated with a Delphix Engine then the log level applies to all the dSources.

Log level description:

- `level=INFO`: This level will print only informational logs. This is the default log level.
- `level=DEBUG`: This level will print informational logs and debug statements.

Setting Maximum Size of Active Log File (`logFileSize`)

Delphix users can set this parameter under the `plugin_logging_parameters` header of the `db2_plugin.conf` file to set the maximum size of the active log file in MB. Once this limit is reached, the plugin will rotate the log as per the retention property defined below. The minimum value of this parameter is 1 MB and maximum value is 10 MB. This parameter only takes in a positive integer value. The default value is 1 MB.

Setting Log Retention Levels (`retention`)

Delphix users can set a retention level for diag logs (`<Db name>.diag.log`) using parameter ***retention*** under the `plugin_logging_parameters` header. As per this parameter, the log files are moved (archived), renamed or deleted once they reach the value set in `logFileSize` parameter. New incoming log data is directed into a new fresh file (at the same location).

By default, this value will be set to a minimum value of 2. The user can change this value and set its value within the range 2 and 50. For example; if ***retention*** is set to 4, the plugin will have the following log files: `<DB Name>.diag.log`, `<DB Name>.diag.log.1`, `<DB Name>.diag.log.2`, `<DB Name>.diag.log.3`, `<DB Name>.diag.log.4`.

- File `<DB Name>.diag.log` is the active log file.
- File `<DB Name>.diag.log.1` is the most recent archive log file
- File `<DB Name>.diag.log.4` is the oldest one

db2_plugin.conf

```
#
# The parameters must not be left padded with spaces. The plugin will otherwise not
# accept the specified parameter
# values and will use the default values.
```

```
#

[plugin_custom_parameters]
#
# If the user has not shared a common group between primary environment user and
instance user.
# Then the user needs to set parameter usersHaveCommonGroup as False.
#

usersHaveCommonGroup=True

#
# During provision operations the plugin checks if the target instance contains a
database name which is identical to
# the source DB name from which the VDB is to be created. By default the plugin will
never allow creating a VDB on the
# instance where the source database already exists, or on any instance that has a
database with the same name. If the
# users still want to create a VDB on an instance which contains a database with the
same name as the source DB then
# they need to set parameter allowSourceDbOnSameInstance as True.
#

allowSourceDbOnSameInstance=False

#
# If the users want to control the number of parallel restores then they can tune the
parameter restorePipelineLimit.
# By default the value of this parameter would be 10.
#

restorePipelineLimit=10

[plugin_logging_parameters]
#
# This flag will set the debug level of plugin logs on the remote server (both
staging and target). There are two levels:
# Info
# Debug
# The above are the only valid values that can be assigned. If any other value is
assigned, the plugin will set the
# default level as Info. the string is case insensitive so info, Info and INFO are
acceptable.
#

level=INFO

#
# This parameter will set the maximum size of the active log file in MB. Once this
limit is reached, the plugin will
# rotate the log as per the retention property defined below. The minimum value of
this parameter is 1 MB and maximum
```

```
# value is 10 MB. This parameter only takes in a positive integer value. The default
value is 1 MB
# logFileSize=<positive integer>
#
logFileSize=1

#
# Whenever the size of <DB Name>.diag.log file exceeds the value provided by
logFileSize parameter then the plugin will
# rename the active log file to <DB Name>.diag.log.<number> and a new log file with
name <DB Name>.diag.log will be
# generated. For example, if LogRetention is set to 4, the plugin will have the
following log files: <DB Name>.diag.log,
# <DB Name>.diag.log.1, <DB Name>.diag.log.2, <DB Name>.diag.log.3, <DB
Name>.diag.log.4.
# File <DB Name>.diag.log.4 will be the oldest one.
# File <DB Name>.diag.log.1 will be the most recent archive log file.
# File <DB Name>.diag.log will be the active log file.
# The minimum and default value for retention flag is 2.
# retention=<positive integer>
#
retention=2
```

Related Topics

- [Db2 Plugin Installation and Upgrade](#)
- [Db2 Support and Requirements](#)

PostgreSQL Environments and Data Sources

This section covers the following topics:


- [Delphix Architecture with PostgreSQL](#)
- [Quick Start Guide for PostgreSQL](#)
- [PostgreSQL Plugin Installation](#)
- [PostgreSQL Support and Requirements](#)
- [Managing PostgreSQL Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with PostgreSQL](#)
- [Provisioning and Managing VDBs from PostgreSQL dSources](#)
- [PostgreSQL Hook Operation Notes](#)
- [Managing Logs Retention](#)

Delphix Architecture with PostgreSQL


Primer: PostgreSQL Replication Methodology

Delphix achieves virtualization for the PostgreSQL database through leveraging PostgreSQL's streaming replication protocol between a Primary (parent) and Secondary (child) database configuration. Delphix uses the secondary, hot-standby database as a source database to ingest from.

In a typical PostgreSQL replication workflow, PostgreSQL leverages a parent and child configuration to maintain high-availability with Write Ahead Logs (WAL - a log file where all the modifications to the database are written before they're applied/written to data files) to maintain sync between both databases. This can be seen in the architectural diagram below.



Delphix leverages this feature by capturing the replication stream via file-based log shipping or Streaming WAL records depending on your PostgreSQL configuration. Below is the same PostgreSQL replication architecture diagram with Delphix use case overlaid:



Core Concepts:

Write Ahead Logs - WAL:

- These are log files that record all modifications to the database before they're applied/written to data files.

File-Based Log Shipping:

- A method of syncing parent and child databases by applying WAL log files.
- One WAL log file can contain up to 16MB of data
- The WAL file is shipped only after it reaches that threshold
- Note: This will cause a delay in replication and also increase the chances of losing data if the parent crashes and logs are not archived.

Streaming WAL Records:

- WAL record chunks are streamed by database servers to keep data in sync.
- The standby server connects to the parent to receive the WAL chunks.
- The WAL records are streamed as they are generated.
- The streaming of WAL records need not wait for the WAL file to be filled.
- *This allows a standby server to stay more up-to-date than is possible with file-based log shipping.*
- By default, streaming replication is asynchronous even though it also supports synchronous replication.

Delphix supports both methods of ingestion (Log Shipping and Streaming WAL Records). It is important to note that both methods have pros/cons. For example, streaming replication doesn't have as much lag between parent and child, as records are sent as they are generated. However, streaming requires both parent and child to be online and able to communicate directly. It also requires the replica to keep up well enough that the parent still has on-disk copies of the WAL the replica needs, and generally requires you to spend extra `pg_xlog` space on retaining extra WAL for the replica.

PostgreSQL Architecture for Streaming Replication

Transaction logs - WAL (Write Ahead Logs) - are replayed from source to the staging environment to maintain data sync between the two.

Steps involved in supporting this architectural setup:

1. The end-user application connecting the PostgreSQL source database may perform read/write queries on the database.

2. The database changes are recorded as WAL segments in the PostgreSQL database under the directory **pg_xlog**.
3. Set up the source PostgreSQL server as Standby Node.
4. Configure replication security by creating a replication user and specifying the authentication protocol.
5. Initiate a base backup on secondary.
6. Configure postgresql.conf file as per the source environment.
7. Start the standby server.

⚠ WAL receiver

The WAL receiver process at the secondary continuously listens for any incoming WAL segments from primary in its receive queue and applies the same on the source database.


⚠ WAL segments

WAL segments are archived when a segment size reaches 16 MB (default).

Delphix Architecture with PostgreSQL

This topic describes the high-level process for adding PostgreSQL environments, linking PostgreSQL data sources to the Delphix Engine, and provisioning virtual databases from PostgreSQL data sources.

Given the preceding primer on Streaming Replication and WAL Logs, A high-level breakdown of the PostgreSQL/ Delphix Architecture breaks down as follows:



To match Delphix Terminology the PostgreSQL-denoted parent Host will be referred to as the Source Host and the PostgreSQL-denoted child Host will be referred to as the Staging Host.

For greater detail - Delphix has three ingestion options that leverage a combination of the PostgreSQL Replication Syncing Method as well as backup ingestion initiation:

1. Delphix Initiated Backup Ingestion with PostgreSQL Streaming Replication.
2. Customer Initiated Backup Ingestion with PostgreSQL Streaming Replication.
3. Customer Initiated Backup Ingestion with WAL log Shipping.

Below is a brief breakdown of the pros/cons of the three ingestions options:

Ingestion Method	Backup Mechanism	Syncing Mechanism	Delphix Considerations
------------------	------------------	-------------------	------------------------

Delphix Backup + Streaming Replication	The initial backup is taken by Plugin through the <code>pg_basebackup</code> method.	Streaming ensures that source and staging databases are closely synched	This method ensures that Delphix has the capability to ingest as frequently as WAL logs are generated by the Source. This leads to less snapshots than a customer-initiated backup, but those snapshots are aligned to policy.
Customer Backup + Streaming Replication	The initial backup is provided by the customer. Backup is taken through the <code>pg_basebackup</code> method.	Streaming ensures that source and staging databases are closely synched	This method ensures that the Delphix Engine dSource is always closely aligned to the Source Database
Customer Backup + Log Shipping	The initial backup is provided by the customer. Backup is taken through the <code>pg_basebackup</code> method.	The less bandwidth-intensive method which does not require source and staging databases to be directly connected at all times	This method does not require all resources to be concurrent, which leads to a less resource-intensive setup.

Environment Setup

Linking Architecture between Postgres and Delphix Engine

At a high level, the Delphix Engine maintains a logical representation of the source database files, from which one can provision virtual databases (VDBs). In order to link a data source and provision a VDB, the following types of environments are required:

Postgres Provisioning Architecture - Source Environment - to - Delphix Engine

A source environment is where the un-virtualized source database runs. The Delphix Engine uses the backup, restore, and replication features of the PostgreSQL DBMS to maintain its internal representation of the source database, to be used for provisioning VDBs. The Delphix Engine must be able to connect to the source environment in order to discover each running source database and to orchestrate the backup, restore, and replication functionality necessary to keep its representation synchronized with the source database. The Delphix Engine is designed to have a minimal impact on the performance of the source database and the source environment.

Target Environments for Postgres

A target environment is where virtualized databases run. PostgreSQL target environments serve two purposes:

1. Since PostgreSQL does not provide a native incremental backup API, a warm standby server (in other words, one in log-shipping mode) must be created where all database files are stored remotely on a Delphix Engine. We refer to the process of applying those WAL files to the staging database as validated sync. During validated sync, we retrieve shipped WAL logs from the source, roll the staging database forward, and create a snapshot of that data state.

2. Once a source database has been set up, you can provision virtual databases from any of the discrete snapshots along the Timeflow mentioned above to any compatible target environment (for more information, see [Requirements for PostgreSQL Target Hosts and Databases](#)). Database files are exported over the network to the target environment, where the virtual database instance runs.

Workflow for PostgreSQL Environments

Linking Environments

Prior to linking a data source, both the source environment and a compatible target environment (to be used for the source database mentioned above) must be added to the Delphix Engine. Prior to provisioning a virtual database, a target environment (with similar settings/versioning to the source environment) must be added to the Delphix Engine.

TARGET HOSTS for PostgreSQL

Container for VDBs

This topic describes the basic concepts involved with provisioning VDBs from Postgres dSources or even other Postgres VDBs.

Once an environment is added to the Delphix Engine, environment discovery takes place. Environment discovery is the process of enumerating PostgreSQL installations and configurations when a source or target environment is added to the Delphix Engine. The discovery process is repeated during environment refresh in order to detect new PostgreSQL installations and clusters.

A dSource is a copy of a physical database that is created when the Delphix Engine links to and loads the database. The Delphix Engine keeps the dSource in sync with the source database in order to facilitate the provisioning of Virtual Databases (VDBs) from the dSource's TimeFlow. While creating a dsource, the Delphix Engine initiates a full database backup of the source database by running `pg_basebackup` on the source host. The initial snapshot of the dSource is derived from this backup.

After obtaining the initial snapshot and linking the dSource, the Delphix Engine keeps the source and staging database in sync by monitoring the source database for new transaction logs on the source host and then applying those transaction logs on the source database. PostgreSQL streaming replication protocol is used to achieve data replication between the source database and the staging database.

When provisioning a VDB, Delphix creates the database with the default Postgres database options. If the database options have been altered on the source database and you wish for the VDB to reflect these same options, they would need to be altered via a Post-Script or by hook operations.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot.

Related Topics

- [PostgreSQL Support and Requirements](#)

Quick Start Guide for PostgreSQL

This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with PostgreSQL database objects in the Delphix Engine. It does not cover any advanced configuration options or best practices, which can have a significant impact on performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix functionality. It assumes you will use the VMware Hypervisor.

Overview

In this guide, we will walk through deploying a Delphix Engine, starting with configuring Source and Target database environment. We will then create a dSource, and provision a VDB.

For purposes of the QuickStart, you can ignore any references to Replication or Masking, such as the engines shown in the diagram below.

Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to the Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online. This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
                    example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
                    the primary interface. Setting this value
                    to "true" will cause all other properties (address,
                    hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
                    "1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
                    for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
                    -- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false

```

```
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24
```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

8. Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`)

9. Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

10. Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the browser at:
```

```
http://<address>
```

Type "exit" to disconnect, or any other commands to **continue** using the CLI.

11. Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
12. Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Virtualization Setup

- Welcome
- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

Welcome

Choose engine type to setup:

Virtualization
 Masking

This wizard will step you through the setup. During this process you will complete the following:

- Create your password for the default "sysadmin" user
- Set the system time
- Configure network and services
- Configure the storage pool
- Configure proxies, SMTP, and LDAP (these are optional)
- Register your software

After setup is complete, you will have two administrators defined:

- The system administrator, "**sysadmin**" with the password you defined. This will be the system administrator for the instance.
- The engine administrator, "**admin**" with the password you defined. This is typically a DBA who will administer all the data managed by the instance.

When setup is complete, log in as engine administrator to begin using your engine.

The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

1. Connect to the Delphix Engine at `http://<Delphix Engine>/login/index.html#serverSetup`.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

⚠ The default domain user created on Delphix Engines from 5.3.1 is known as **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications

Choose your option to set up system time in this section. For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click **Next** to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication Service

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:

- a. Manually highlight the registration code and copy it to clipboard. Or,
- b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.
4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

 Although your Delphix Engine will work without registration, we strongly recommend that you register each engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

Requirements for PostgreSQL Source Databases

Source Database Requirements

- The Delphix source database can be in READ-WRITE or READ only mode. It can be a standby database(streaming site) as well.
- The Staging host must have access to a PostgreSQL role on the source side that has replication, and login privileges. This can be the built-in PostgreSQL role or a newly-created role (for example, delphix).

To create a **new** role **for** use with the Delphix Engine, use the following command:

```
SQL> CREATE USER delphix WITH REPLICATION ENCRYPTED [ PASSWORD 'password'];
```

- You must make the following changes to postgresql.conf (for more information, see the [Server Configuration](#) chapter in the PostgreSQL documentation):
- TCP/IP connectivity must be configured to allow the role mentioned above to connect to the source database from the Delphix Engine and from the standby DBMS instance set up by the Delphix Engine on the staging environment. This can be done by modifying the parameterlisten_addresses, which specifies the TCP/IP addresses on which the DBMS is to listen for connections from client applications.

```
listen_addresses = '*'           # defaults to 'localhost'; use '*' for all
```

- The value of `max_wal_senders`, which specifies the maximum number of concurrent connections from standby servers or streaming base backup clients, must be increased from its desired value by four. That is, in addition to the allowance of connections for consumers other than the Delphix Engine, there must be an allowance for four additional connections from consumers set up by the Delphix Engine.

The **default** value `max_wal_senders` is zero, meaning replication is disabled. In **this** configuration, the value of `max_wal_senders` must be increased to two **for** the Delphix Engine:

```
max_wal_senders = 4           # Default is 0
```

- The value of `wal_level` which determines how much information is written to the write-ahead log (WAL), must be set to `archive` or `hot_standby` to allow connections from standby servers. The value `logical_wal_level` (introduced in PostgreSQL 9.4) is also supported.

The **default** value of `wal_level` is `minimal`, which writes only the information needed to recover from a crash or immediate shutdown to the WAL archives. In **this** configuration, you must add the logging required **for** WAL archiving as follows:

```
wal_level = archive           # Default is minimal
```

- In releases prior to PostgreSQL 9.6.x, parameter `'wal_level'` allowed the values `'archive'` and `'hot_standby'`. These values are still accepted but mapped to `'replica'`. For example, The plugin will work with `ALTER SYSTEM SET wal_level TO 'replica'` and `ALTER SYSTEM SET wal_level TO 'hot_standby'`.
- You must configure PostgreSQL to allow PostgreSQL client connections and replication client connections from the staging target environment. To configure appropriately, add the following entries to `pg_hba.conf`:

```
host all <role> <ip-address_of_delphix_engine>/32 <auth-method> host all <role>
<ip-address_of_staging_target>/32 <auth-method> host replication <role> <ip-
address_of_staging_target>/32 <auth-method>
```

- `<auth-method>` can be set to `trust`, `scram-sha-256`, or `md5` on source. Delphix inherits the same authentication method from the source while dSource creation and VDB provisioning. For more information on how to configure, `pg_hba.conf` see the [Client Authentication](#) chapter in the PostgreSQL documentation.
- `wal_keep_segments` parameter in `postgresql.conf` file on the Source environment should be large enough to support the WAL sync process once the dSource is building up. Postgres 13 onwards, the configuration parameter `wal_keep_segments` is changed to `wal_keep_size`. It is specified in megabytes rather than the number of files as with the `wal_keep_segments` parameter. If you previously used `wal_keep_segments`, the following formula will give you an approximate equivalent setting:
`wal_keep_size = wal_keep_segments * wal_segment_size` (typically 16MB).

- It is mandatory to have "port" (can be commented or not commented) in postgresql.conf file.

Requirements for PostgreSQL Target Hosts and Databases

Target Environment Requirements

1. The operating system and architecture of the target environment must match those of the source environment. It is recommended that the source and the target environments should be identical and hardware configurations should match.
 2. For supported OS version and DBMS version, see [PostgreSQL Matrix](#). The underlying Operating System for both the Source and Staging environment should be amongst these two versions.
 3. There must be an installation of PostgreSQL on the target environment that is compatible with the installation of PostgreSQL on the source environment. Two installations of PostgreSQL are compatible if and only if:
 - a. They share the same vendor (for example, PostgreSQL is incompatible with EnterpriseDB Postgres Plus Advanced Server).
 - b. They share the same major version number (for example, 9.5.4 is compatible with 9.5, 9.5.3; however, it is incompatible with 9.3, 9.3.24, or 9.2).
 - c. They are compiled against the same architecture (in other words, 32-bit and 64-bit installations of Postgres are incompatible).
 - d. They are compiled with the same WAL segment size. The default WAL segment size of 16 MB is rarely changed in practice, so almost all installations of PostgreSQL are compatible with each other in terms of WAL segment size.
1. There must be an operating system user (e.g. PostgreSQL) with the following privileges:
 - a. The Delphix Engine must be able to make an SSH connection to the target environment using the operating system user.
 - b. The operating system user must have read and execute privileges on the PostgreSQL binaries installed on the target environment.
 - c. The operating system user must have permission to run mount and umount as the superuser via sudo with neither a password nor a TTY. See [Sudo Privilege Requirements for PostgreSQL Environments](#) for examples of the /etc/sudoers file on different operating systems.
 1. There must be a directory on the target environment where the Delphix Engine Plugin can be installed (for example, /var/tmp) with the following properties:
 - a. The Plugin directory must be writable by the operating system user mentioned above.
 - b. The Plugin directory must have at least 1.5 GB of available storage.
 1. There must be a mount point directory (for example, /mnt) that will be used as the base for mount points that are created when provisioning a VDB with the following properties:
 - a. The mount point directory must be writable by the operating system user mentioned above.
 - b. The mount point directory should be empty.
 1. TCP/IP connectivity to and from the source environment must be configured as described in [General Network and Connectivity Requirements](#).
 2. Hostname and IP must be correctly set in /etc/hosts file, for example [postgres@source ~]\$ hostname -i
The output of "hostname -i" command should produce the correct result as the IP address of the server. For example:
[postgres@source postgres]\$ hostname -i
10.110.207.113
 - When using the External Backup method for ingestion the directories being used for keeping the Backup file and WAL file should be accessible by OS User.
 - Zip must be installed on the Staging/Target Host for External Backup Support.

- The backup command and File should be in the below format:
PostgreSQL_T<YYYYMMDD>.zip where YYYY denotes Year, MM denotes month, DD denotes date. For example PostgreSQL_T20190314.zip.
The above zip file should contain a Database Full Backup from the Source created in tar format. The below flags can be used with pg_basebackup for taking Source Database backup:

For PostgreSQL versions 9.4.x, 9.5.x, and 9.6.x

```
<PATH_TO_PG_BIN>/pg_basebackup -p<PORT>-D<EMPTY_BACKUP_DIR>-F t -v -w -P -x
```

For PostgreSQL versions 10.x and 11.x

```
<PATH_TO_PG_BIN>/pg_basebackup -p <PORT> -D <EMPTY_BACKUP_DIR> -F t -v -w -P -Xs
```

6. If there's no PostgreSQL instance running on the Target host, however, we have a PostgreSQL installation, even then the plugin relies on the "DELPHIX_PG_PATH" variable to discover the environment. In the absence of "DELPHIX_PG_PATH" variable or if the value of "DELPHIX_PG_PATH" variable is NULL then Linux "find" command will be used for the Environment discovery which may impact the overall performance. Hence, it is preferred to create this Environment Variable with correct entries. In order to optimize the performance, it is preferred to create an Environment Variable "DELPHIX_PG_PATH" which should be accessible by OS users.
Below should be the syntax for "DELPHIX_PG_PATH" variable:
DELPHIX_PG_PATH="binary_path:data_path1;binary_path:data_path2;"
For example:
DELPHIX_PG_PATH="/usr/pgsql-9.6/bin:/var/lib/pgsql/9.6/data;/opt/edb/as9.6/bin:/opt/edb/as9.6/data;/usr/pgsql-9.6/bin:/tmp/TESTING/data;"

 The variable "**DELPHIX_PG_PATH**" must be available to the environment user in a non-interactive way. You can test this variable for non-interactive logins using `ssh your_username@v-cardtr.tmydb.kv.aval "env | grep DELPHIX_PG_PATH"`.

Adding a PostgreSQL Environment

Prerequisites

- Make sure that the staging environment in question meets the requirements described in [Requirements for PostgreSQL Hosts and Databases](#)

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next, to Environments, click the Actions (...) menu and select **Add Environment**.
5. In the Add Environment dialog, select **Unix/Linux**.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter Name for the Environment.
9. Enter the **Host IP address** or **hostname**.
10. Enter the **SSH port**. The default value is 22.
11. Enter an **OS Username** for the Environment which is currently 'postgres'.
12. Select **Login Type**.
— Username and Password - enter the OS username and password

- Username and Public Key - enter the OS username.
- Password Vault - select from an existing Enterprise Password Vault

i Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For Password, enter the password associated with the user in step 11.
14. If you want to use Public Key Encryption for logging into your environment:

a. Select Public Key for the Login Type.

b. Click View Public Key.

c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.

i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).

ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.

iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.


The public key needs to be added only once per user and per environment.

You can also add public key authentication to an environment user's profile by using the command-line interface, as explained in the topic [CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users](#).

15. For Password Login, click Verify Credentials to test the username and password.
16. Enter **Toolkit Path** (make sure Toolkit path does not have spaces).
17. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
18. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
19. Click **Submit**.

As the new environment is added, you will see two jobs running in the Delphix platform Job History, one to Create and Discover an environment, and another to create an environment. When the jobs are complete, you will see the new environment added to the list in the Environments tab. If you do not see it, click the Refresh icon in your browser.

18. Once the environment is discovered, further linking would require adding a source config to the above-discovered installation. Please refer to [Linking a PostgreSQL dSource](#) for more information.



Linking a PostgreSQL Data Source

Prerequisites

- The source and staging instances must meet the host requirements
 - Databases must meet container requirements
- For more information refer to [Requirements for PostgreSQL Source Hosts and Databases](#)

Limitation

It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.

Procedure

1. Log in to the **Delphix Management** application.
2. Select **Manage > Environments**.
3. From the **Databases** tab, select a repository from which you want to create your dSource and click the **+** icon.
4. In the **Add Database** dialog window enter the **Name** for your source config and click **Add**.
5. Select your source config and click the **Add dSource** link located to the right.
6. In the Source tab click on the **+Add** icon located next to **Delphix Initiated Backup** Parameter.

⚠ The Add Button provided for Delphix Initiated Backup/External Backup - Streaming Replication is clickable multiple times but should be used **ONLY** once since the solution is built only to support a single backup path. Providing Multiple Paths will error out the linking process.

The following fields are mandatory for this process :

- **Delphix Initiated Backup Flag** - Checkbox should be selected.
- **PostgresDB Replication User**
- **PostgresDB Replication User Password**
- **Source Host Address**

- **Source Instance Port Number**
- **Staging Instance Port Number** - must be unique for a given instance

Data Type
AppData

External Backup
dSource Creation through external backup

+ Add

Allow the dSource to be in sync through streaming replication if WalLogs are not provided.
Allow the dSource to be in sync with Source through streaming replication if WalLogs are not provided. "Delphix Initiated Backup" parameters are mandatory for this.

Delphix Initiated Backup Flag
Allow Delphix to initiate the backups from source database. "Delphix Initiated Backup" parameters are mandatory for ingestion.

Delphix Initiated Backup/External Backup - Streaming Replication
List of parameters for Delphix initiated backup and external backup with streaming replication

PostgresDB Replication User *
delphix Delete

Replication User for postgres database

PostgresDB Replication User Password *

Replication password for postgres database

Source Host Address *
centos78pg12src.dlpxdc.co

Source Host Address

Source Instance Port Number *
5432

Port number for postgres source database

+ Add

Staging Instance Port Number *
5433

Port number for postgres staging database

Config Settings
Custom Database-Level config settings

+ Add

Cancel Back **Next** Submit

7. Optionally, database configuration parameters can be defined during the linking operation in the Config Settings section.

Config Settings

Custom Database-Level config settings

Delete

Comment Property

Select this option to comment out the provided property name in the configuration file

+ Add

i After dSource creation, users can now configure the parameters and their values in the `postgresql.conf` file through the "Config Settings" section on the Delphix Engine UI. For example, to support **Replication Slots** through UI, users can now provide the "primary_slot_name" parameter and its value through the UI. This will update the already existing parameter value in the `postgresql.conf` file with the value provided by the user.

Similarly, if a user wants to disable a parameter then they can provide the parameter name and select the check box **Comment Property**. This will comment out the parameter in the `postgresql.conf` configuration file.

Click **Next**.

8. In the **dSource Configuration** tab, enter a dSource name and select a group for your dSource. Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
9. In the **Data Management** tab specify your **Staging Environment**, **User**, and **Mount Base**, then click **Next**.
10. If required apply policy details to the dSource, then click **Next**.
11. Users can provide bash scripts as Plugin defined hooks. These hooks are executed after data ingestion but before snapshot operation, click **Next**.
12. Review the dSource Configuration and Data Management information in the **Summary** section.
13. Click **Submit**.

The Delphix Engine initiates two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking **Active Jobs** from the top menu bar, or by selecting **System > Event Viewer**. When the jobs have successfully completed, the database icon will change to a dSource icon on the **Environments > Host > Databases** screen, and the dSource will also appear in the list of Datasets under its assigned group.

⚠ The dSource Configuration Screen

After you have created a dSource, the dSource Configuration tab allows you to view information about it and make modifications to its policies and permissions. In the Datasets panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the Source files, Data

Management configuration, and Hook Operations. For more information, refer to [Data Management Settings for PostgreSQL Data Sources](#).

Provisioning a PostgreSQL VDB

Prerequisites

- You will need to have linked a dSource from a staging instance, as described in [Linking a PostgreSQL dSource](#) or have created a VDB from which you want to provision another VDB
- You should have set up the POSTGRES target environment with necessary requirements as described in [PostgreSQL Support and Requirements](#)
- Make sure you have the required Instance Owner permissions on the target instance and environment
- The method for [Database Permissions for Provisioned PostgreSQL VDBs](#) is decided before the provisioning

Procedure

1. Navigate to **Manage**, and select **Datasets**.
2. Select a dSource and a snapshot from which you want to provision. Click the provision VDB icon to open the **provision VDB** wizard.
3. Select a target environment from the left pane, and an Installation to use from the dropdown list of available PostgreSQL instances on that environment.
4. Set the **Environment User** to be the Instance Owner. **Note:** The picking of instance owner is only possible if you have multiple environment users set on that host.
5. You will see the **Target Configuration** section where you need to specify **Mount Path**.

The screenshot shows the 'Provision VDB' wizard interface. On the left, a vertical progress bar lists the steps: Preparation, Source, Provision Point, Target Environment, **Target Configuration** (highlighted with a blue dot), Configuration, Policies, Masking, Hooks, and Summary. The main content area is titled 'Target Configuration' and contains the instruction 'Configure the target environment.' Below this, there is a 'Mount Path' label and a text input field containing the path '/mnt/provision/V/mnt/provision/r82d115a_bfvyMTVs'. At the bottom of this section, there is a checkbox labeled 'Additional Mount Points' which is currently unchecked.

6. Optionally, set the database configuration parameters for the VDB. Users can disable a configuration parameter by selecting the check box Comment Property. This will comment out the parameter in the `postgresql.conf` configuration file. Click **Next**.
7. On the configuration page enter the PostgreSQL database name as the vFiles Name value.

8. Select a **Target Group** for the VDB and click the green **Plus** icon to add a new group, if necessary.
9. Select a **Snapshot Policy** for the VDB then click **Next**.
10. Specify any desired hook operations.
11. Review the Provisioning Configuration and Data Management information.
12. Click **Submit**.

Once the VDB provisioning has successfully completed, if the source and target instance ids are not the same, you may want to grant secadm and dbadm on the database to the target instance id. Please refer to [Database Permissions for Provisioned PostgreSQL VDBs](#) for more information.

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.
2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test the sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision new VDBs from that masked VDB to quickly provide safe data to development and QA teams.

Related Topics

- [PostgreSQL Support and Requirements](#)
- [Managing PostgreSQL Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with PostgreSQL](#)
- [Provisioning and Managing VDBs from PostgreSQL dSources](#)

PostgreSQL Plugin Installation

The Plugin is provided as a zip file in the Delphix download site. Follow below steps to download plugin:

1. In the web browser, go to the Delphix [download](#) site.
2. Login to the download site using email and password credentials.
3. Navigate to the required version number of Delphix Engine.
4. Go to the Plugins > PostgreSQL > Plugin_<version_number>.zip folder and download the zip file.

Refer to the section [Delphix Engine Plugin Management](#) for further details.

Related Links

- [PostgreSQL Support and Requirements](#)
- [Delphix Engine Plugin Management](#)

PostgreSQL Support and Requirements

In order to begin using PostgreSQL environments with Delphix, you will need to configure the source and target hosts with the requirements described in this section.

This section covers the following topics:

- [PostgreSQL Compatibility Matrix](#)
- [Requirements for PostgreSQL Hosts and Databases](#)
- [Network and Connectivity Requirements for PostgreSQL Environments](#)
- [Sudo Privilege Requirements for PostgreSQL Environments](#)
- [Sudo File Configuration Examples for PostgreSQL Environments](#)

PostgreSQL Compatibility Matrix

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Key:

Color	Supported?
	Yes
	No

RHEL/CentOS/SUSE

Supported OS Version	Supported DBMS Version							
	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	Enterprise DB (Open-source & EDB)
RHEL/CentOS 7.3	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.4	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.5	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.6	✓	✓	✓	✓	✓	✓	✗	
RHEL/CentOS 7.7	✓	✓	✓	✓	✓	✓	✗	

	Supported DBMS Version							
Supported OS Version	9.4.x (Open Source & EDB)	9.5.x (Open Source & EDB)	9.6.x (Open-source & EDB)	10.x (Open-source & EDB)	11.x (Open-source & EDB)	12.x (Open-source & EDB)	13.x (Open-source & EDB)	EnterpriseDB (Open-source & EDB)
RHEL/CentOS 7.8	✗	✗	✓	✗	✗	✓	✗	
RHEL/CentOS 7.9	✗	✗	✓	✓	✗	✓	✓	■
RHEL/CentOS 8.3	✗	✗	✗	✗	✗	✓	✓	■
SUSE 11	✓	✗	✗	✗	✗	✗	✗	
SUSE 12	✓	✓	✓	✓	✓	✗	✗	

where **x** represents the minor version.

- ⚠ All the available PostgreSQL Database minor versions will be supported for the above-mentioned major versions.
- Make sure to install libncurses.5 or libncurses.6 in the host, for supporting RHEL and SUSE. You can refer [KBA 5622](#) for actionable steps.

Plugin/Delphix Engine Compatibility

Plugins should be installed on compatible Delphix Engines per the table below:

- ⚠ The plugin versions from 1.0.1 till 1.3.2 are currently in extended support. Please reach out to Delphix customer support for any queries.

- ⚠ Currently, PostgreSQL Plugin 1.4.1 is not supported on SUSE 12.x.

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.0.x	✓	✓	✓	✓	✓	✓	✓
6.0.1.x	✓	✓	✓	✓	✓	✓	✓
6.0.2.x	✓	✓	✓	✓	✓	✓	✓
6.0.3.x	✓	✓	✓	✓	✓	✓	✓
6.0.4.x	✓	✓	✓	✓	✓	✓	✓
6.0.5.x	✓	✓	✓	✓	✓	✓	✓
6.0.6.x	✓	✓	✓	✓	✓	✓	✓
6.0.7.x	✓	✓	✓	✓	✓	✓	✓
6.0.8.x	✓	✓	✓	✓	✓	✓	✓
6.0.9.x	✓	✓	✓	✓	✓	✓	✓

Delphix Engine Version	Plugin Version						
	1.4.0	1.4.1	1.4.2	1.4.3	1.4.4	1.5.0	1.5.1
6.0.10.x	✓	✓	✓	✓	✓	✓	✓
6.0.11.x	✓	✓	✓	✓	✓	✓	✓
6.0.12.x	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓

Unsupported PostgreSQL Versions and Features

- PostgreSQL versions below 9.4 i.e (9.2, 9.3) are NOT supported by the plugin.
- It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.
- Broadly there are two features that are not supported by the plugin:
 - a. Unlogged Tables: All the tables which are not logged (unlogged) will not be supported by the Plugin.
 - b. Point In Time recovery: Currently, the Plugin doesn't support Point in Time recovery of the PostgreSQL database.

Related Topics

- [PostgreSQL Support and Requirements](#)

Requirements for PostgreSQL Hosts and Databases

In order to begin using PostgreSQL environments with Delphix, you will need to configure the source and target with the requirements described on this page.

PostgreSQL Hosts and Databases

On each host with Postgres, there must be an operating system user configured to the required specifications for Delphix, as explained in the table below. These requirements apply to both source and target environments. However, target environments have additional requirements which are detailed in the ‘Target Host Requirements’ section below.

 The PostgreSQL plugin requires the use of the default user which gets created during the Postgres DB installation process. Usually, 'postgres'. Please add the environment using this user and create all the datasets with the same.

Source Database Requirements

Database Requirement	Explanation
<p>The source database can be in READ-WRITE or READ only mode. It can be a standby database(streaming site) as well.</p>	
<p>The Staging host must have access to a PostgreSQL role on the source side that has replication, and login privileges. This can be the built-in PostgreSQL role or a newly-created role (for example, delphix).</p>	<p>Creating a Role for Use with the Delphix Engine</p> <p>To create a new role for use with the Delphix Engine, use the following command:</p> <pre>SQL> CREATE USER delphix WITH REPLICATION ENCRYPTED [PASSWORD 'password'];</pre>
<p>You must make the following changes to postgresql.conf:</p> <p>modifying the parameter <code>listen_addresses</code>, which specifies the TCP/IP addresses on which the DBMS is to listen for connections from client applications.</p>	<p>Note: TCP/IP connectivity must be configured to allow the role mentioned above to connect to the source database from the Delphix Engine and from the standby DBMS instance set up by the Delphix Engine on the staging environment.</p> <p>listen_addresses configuration</p> <p>The simplest way to configure PostgreSQL is so that it listens on all available IP interfaces:</p> <pre>listen_addresses = '*' # Default is 'localhost'</pre> <p>(for more information, see the Server Configuration chapter in the PostgreSQL documentation)</p>

Database Requirement	Explanation
<p>The value of <code>max_wal_senders</code> , which specifies the maximum number of concurrent connections from standby servers or streaming base backup clients, must be increased from its desired value by two. That is, in addition to the allowance of connections for consumers other than the Delphix Engine, there must be an allowance for two additional connections from consumers set up by the Delphix Engine.</p>	<p>max_wal_senders Configuration</p> <p>The default value <code>max_wal_senders</code> is zero, meaning replication is disabled. In this configuration, the value of <code>max_wal_senders</code> must be increased to two for the Delphix Engine:</p> <pre>max_wal_senders = 2 # Default is 0</pre>
<p>The value of, <code>wal_level</code> which determines how much information is written to the write-ahead log (WAL), must be set to <code>archive</code> or <code>hot_standby</code> to allow connections from standby servers. The value "logical" for <code>wal_level</code> (introduced in PostgreSQL 9.4) is also supported.</p> <p>In releases prior to PostgreSQL 9.6.x, parameter 'wal_level' allowed the values 'archive' and 'hot_standby'. These values are still accepted but mapped to 'replica'. For example, The plugin will work with <code>ALTER SYSTEM SET wal_level TO 'replica'</code> and <code>ALTER SYSTEM SET wal_level TO 'hot_standby'</code>.</p>	<p>wal_level Configuration</p> <p>The default value of <code>wal_level</code> is <code>minimal</code>, which writes only the information needed to recover from a crash or immediate shutdown to the WAL archives. In this configuration, you must add the logging required for WAL archiving as follows:</p> <pre>wal_level = archive # Default is minimal</pre>
<p>You must configure PostgreSQL to allow PostgreSQL client connections and replication client connections from the staging target environment.</p>	<p>To configure appropriately, add the following entries to <code>pg_hba.conf</code>:</p> <pre>pg_hba.conf Configuration host all <role> <ip-address_of_delphix_engine>/ 32 <auth-method> host all <role> <ip-address_of_staging_target>/ 32 <auth-method> host replication <role> <ip- address_of_staging_target>/32 <auth-method></pre>

Database Requirement	Explanation
<p><auth-method> can be set to trust, scram-sha-256, or md5 on source. Delphix inherits the same authentication method from the source while dSource creation and VDB provisioning. For more information on how to configure, pg_hba.conf see the Client Authentication section in the PostgreSQL documentation.</p>	
<p><code>wal_keep_segments</code> parameter in postgresql.conf file on the Source environment should be large enough to support the WAL sync process once the dSource is building up.</p> <p>Postgres 13 onwards, the configuration parameter <code>wal_keep_segments</code> is changed to <code>wal_keep_size</code>. It is specified in megabytes rather than the number of files as with the <code>wal_keep_segments</code> parameter. If you previously used <code>wal_keep_segments</code>, the following formula will give you an approximate equivalent setting:</p> $\text{wal_keep_size} = \text{wal_keep_segments} * \text{wal_segment_size} \text{ (typically 16MB).}$	
<p>It is mandatory to have "port" (can be commented or not commented) in postgresql.conf file.</p>	

Target Host Requirements

Host Requirement	Explanation
<p>The operating system and architecture of the target environment must match those of the source environment.</p>	<p>It is recommended that the source and the target environments should be identical and hardware configurations should match.</p>

Host Requirement	Explanation
<p>Compatible Operating System supported by the Plugin is CentOS 7.3/7.4/7.5/7.6/7.7/7.8/7.9 and RHEL 7.3/7.4/7.5/7.6/7.7/7.8/7.9.</p>	<p>The underlying Operating System for both the Source and Staging environment should be amongst these two versions.</p>
<p>There must be an installation of PostgreSQL on the target environment that is compatible with the installation of PostgreSQL on the source environment.</p>	<p>Two installations of PostgreSQL are compatible if and only if:</p> <ol style="list-style-type: none"> 1. They share the same vendor (for example, PostgreSQL is incompatible with EnterpriseDB Postgres Plus Advanced Server). 2. They share the same major version number (for example, 9.5.4 is compatible with 9.5, 9.5.3; however, it is incompatible with 9.3, 9.3.24, or 9.2). 3. They are compiled against the same architecture (in other words, 32-bit and 64-bit installations of Postgres are incompatible). 4. They are compiled with the same WAL segment size. The default WAL segment size of 16 MB is rarely changed in practice, so almost all installations of PostgreSQL are compatible with each other in terms of WAL segment size.
<p>There must be an operating system user (e.g Postgres) with the following privileges:</p> <ol style="list-style-type: none"> 1. The Delphix Engine must be able to make an SSH connection to the target environment using the operating system user. 2. The operating system user must have read and execute privileges on the PostgreSQL binaries installed on the target environment. 3. The operating system user must have permission to run mount and unmount as the superuser via sudo with neither a password nor a TTY. 	<p>See Sudo Privilege Requirements for PostgreSQL Environments for further explanation of the commands, and Sudo File Configuration Examples for PostgreSQL Environments for examples of the /etc/sudoers file on different operating systems.</p>
<p>There must be a directory on the target environment where the Delphix Engine Plugin can be installed (for example, /var/tmp) with the following properties:</p> <ol style="list-style-type: none"> 1. The Plugin directory must be writable by the operating system user mentioned above. 2. The Plugin directory must have at least 1.5 GB of available storage. 	

Host Requirement	Explanation
<p>There must be a mount point directory (for example, /mnt) that will be used as the base for mount points that are created when provisioning a VDB with the following properties:</p> <ol style="list-style-type: none"> 1. The mount point directory must be writable by the operating system user mentioned above. 2. The mount point directory should be empty. 	
<p>TCP/IP connectivity to and from the source environment must be configured as described in General Network and Connectivity Requirements.</p>	
<p>Hostname and IP must be correctly set in /etc/hosts file, for example [postgres@source ~]\$ hostname -i</p>	<p>The output of “hostname -i” command should produce the correct result as the IP address of the server.</p> <p>For example:</p> <pre>[postgres@source postgres]\$ hostname -i 10.110.207.113</pre>
<p>When using the External Backup method for ingestion the directories being used for keeping Backup file and WAL file should be accessible by OS User.</p>	
<p>Zip must be installed on the Staging/ Target Host for External Backup Support.</p>	
<p>The backup command and File should be in the below format:</p> <p>PostgreSQL_T<YYYYMMDD>.zip where YYYY denotes Year, MM denotes month, DD denotes date. For example, PostgreSQL_T20190314.zip.</p> <p>The above zip file should contain a Database Full Backup from the Source created in tar format.</p>	<p>The below flags can be used with pg_basebackup for taking Source Database backup:</p> <p>For PostgreSQL versions 9.4.x,9.5.x,9.6.x</p> <pre><PATH_TO_PG_BIN>/pg_basebackup -p <PORT> -D <EMPTY_BACKUP_DIR> -F t -v -w -P -x</pre> <p>For PostgreSQL versions 10.x and 11.x</p> <pre><PATH_TO_PG_BIN>/pg_basebackup -p <PORT> -D <EMPTY_BACKUP_DIR> -F t -v -w -P -Xs</pre>

Host Requirement	Explanation
<p>If there's no PostgreSQL instance running on the Staging/Target host, however, we have a PostgreSQL installation, even then the plugin relies on the "DELPHIX_PG_PATH" variable to discover the environment.</p> <p>In the absence of "DELPHIX_PG_PATH" variable or if the value of "DELPHIX_PG_PATH" variable is NULL then Linux "find" command will be used for the Environment discovery which may impact the overall performance. Hence, it is preferred to create this Environment Variable with correct entries.</p> <p>In order to optimize the performance, it is preferred to create an Environment Variable "DELPHIX_PG_PATH" which should be accessible by OS users.</p>	<p>Below should be the syntax for "DELPHIX_PG_PATH" variable: DELPHIX_PG_PATH="binary_path:data_path1;binary_path:data_path2;"</p> <p>For example: DELPHIX_PG_PATH="/usr/pgsql-9.6/bin:/var/lib/pgsql/9.6/data;/opt/edb/as9.6/bin:/opt/edb/as9.6/data;/usr/pgsql-9.6/bin:/tmp/TESTING/data;"</p>

 The variable "**DELPHIX_PG_PATH**" must be available to the environment user in a non-interactive way. You can test this variable for non-interactive logins using `ssh <your_username>@<target_host> "env | grep DELPHIX_PG_PATH"`.

Related Topics

- [PostgreSQL Support and Requirements](#)
- [Sudo Privilege Requirements for PostgreSQL Environments](#)
- [Sudo File Configuration Examples for PostgreSQL Environments](#)
- [Network and Connectivity Requirements for PostgreSQL Environments](#)

Network and Connectivity Requirements for PostgreSQL Environments

Overview

This topic covers the general network and connectivity requirements for the Delphix Engine, including connection requirements, port allocation, and firewall and Intrusion Detection System (IDS) considerations. For platform-specific network and connectivity requirements, see the relevant topics under the Requirements section for each platform.

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine

Protocol	Port Numbers	Use
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI. See Network Performance Tool .

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

SSHD Configuration

Target Unix environments are required to have `sshd` running and configured such that the Delphix Engine can connect over `ssh`.


The Delphix platform expects to maintain long-running, highly performant `ssh` connections with remote Unix environments.

Network and Connectivity Requirements for PostgreSQL

- IP connections must exist between the Delphix Engine and the target environments.
- The Delphix Engine uses an **SSH** connection to each target environment, **NFS** connections from each target environment to the Delphix Engine, and **PostgreSQL client** connections to the virtual databases on the target environment.
- Once connected to a staging target environment through **SSH**, the Delphix Engine initiates a **PostgreSQL replication client** connection from the target environment to the source environment.

Port Allocation for PostgreSQL Environments

The following diagram describes the port allocations for PostgreSQL environments. It illustrates the ports that we recommend to be open from Delphix to remote services, to the Delphix Engine, and to the Target Environments.



Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to the target database environment
TCP	xxx	PostgreSQL client connections to the PostgreSQL instances on the target environments (port 5432 by default)

Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP/UDP	111	Remote Procedure Call (RPC) port mapper used for NFSv3 mounts

Protocol	Port Number	Use
TCP	1110	Network Status Monitor (NSM) client from target hosts to the Delphix Engine
TCP	2049	NFS client from target hosts to the Delphix Engine (NFSv3 and NFSv4)
TCP	54043	Client mount daemon (NFSv3 only)
TCP	54044	Lock state notification service (NFSv3 only)
TCP	54045	Network Lock Manager (NLM) client from target hosts to Delphix Engine (NFSv3 only)
UDP	33434 - 33464	Traceroute from the target database server to the Delphix Engine (optional)

Port Allocation Between Source and Staging Target Environments

Outgoing	Incoming	Protocol	Port Number	Use
Target Environment	Source Environment	PostgreSQL replication client	xxx	PostgreSQL replication client connection to the PostgreSQL instances on the source environment (port 5432 by default)

Related Topics

- [PostgreSQL Support and Requirements](#)

Sudo Privilege Requirements for PostgreSQL Environments

This topic describes the rationale behind specific sudo privilege requirements for virtualizing PostgreSQL Databases.

The sudo configuration exists as /etc/sudoers file.

Below is the example of sudo configuration file contents as mentioned above for Postgres environment.

```
$ vi /etc/sudoers
```

```
Defaults:postgres !requiretty
```

```
postgres ALL=NOPASSWD: /bin/mount,/bin/umount,/bin/mkdir,/bin/rmdir
```

Privilege	Sources	Targets	Rationale
mkdir/rmdir	Not Required	Required	Delphix dynamically creates and removes directories under the provisioning directory during VDB operations.
mount/umount	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because mount and unmount are typically reserved for superuser.

 It is required to specify the NOPASSWD qualifier within the "sudo" configuration file, as shown here: [Sudo File Configuration Examples for PostgreSQL Environments](#). This ensures that the "sudo" command does not demand the entry of a password, even for the "display permissions" (i.e. "sudo -l") command.

AppData Mount Options

Depending on the NFS version used options vers=3 or vers=4.x is added (x varies depending on what that platform supports. e.g. vers=4 or vers=4.1)	
Linux (NFSv3)	-t nfs -o rw,fg,hard,rsize=1048576,wsize=1048576,nointr,ti meo=600,tcp,noacl,vers=3
Linux (NFSv4)	-t nfs4 -o rw,fg,hard,rsize=1048576,wsize=1048576,nointr,ti meo=600,sec=sys,tcp,noacl
<p>Note :</p> <p>(For some flavors of Linux and NFSv4.1, additional optional 'v4.1' is added)</p> <ol style="list-style-type: none"> "port=2049" option is added for all platforms. 	

unmount options

"-f" is used for all platforms. For Linux, "-lf" is used.

Mount and Unmount Options Subject to Change

Please note that the mount and unmount options listed above are subject to change. For example, if Delphix finds that a certain option improves performance, Delphix may add, remove or change options at anytime. Therefore, it is highly recommended to create the sudo profiles using wildcards that allow any number of options.

Related Topics

- [PostgreSQL Support and Requirements](#)

Sudo File Configuration Examples for PostgreSQL Environments

This topic describes the rationale behind specific sudo privilege requirements for virtualizing PostgreSQL Databases.

The sudo configuration exists as `/etc/sudoers` file.

Requiretty settings

Delphix requires that the `requiretty` setting be disabled for all Delphix users with `sudo` privileges.

Configuring `sudo` Access on Linux for PostgreSQL Target Environments

Below is the example of sudo configuration file contents as mentioned above for the Postgres environment.

```
$ vi /etc/sudoers
Defaults:postgres !requiretty
postgres ALL=NOPASSWD: /bin/mount,/bin/umount,/bin/mkdir,/bin/rmdir
```

 Note that the following examples are for illustrative purposes and the sudo file configuration options are subject to change.

Example 1

This example restricts the PostgreSQL user's use of `sudo` privileges to the directory `/postgres`.

Note that wildcards are allowed for the options on `mount` and `umount` because those commands expect a fixed number of arguments after the options. The option wildcard on the `mount` command also makes it possible to specify the file-system being mounted from the Delphix Engine.

However, wildcards are not acceptable on `mkdir` and `rmdir` because they can have any number of arguments after the options. For those commands, you must specify the exact options (`-p` , `-p -m 755`) used by the Delphix Engine.

Example: `/etc/sudoers` File Configuration on the Target Environment for `sudo` Privileges on the VDB Mount Directory Only (Linux OS)

```
Defaults:postgres !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount * /postgres/*, \
/bin/umount * /postgres/*, \
```

```

/bin/umount          /postgres/*, \
/bin/mkdir -p       /postgres/*, \
/bin/mkdir -p -m 755 /postgres/*, \
/bin/mkdir          /postgres/*, \
/bin/rmdir          /postgres/*

```

Example 2

This example restricts the PostgreSQL user's use of `sudo` privileges to the directory `/postgres`, restricts the mount commands to a specific Delphix Engine hostname and IP, and does not allow user-specified options for the `umount` command.

This configuration is more secure, but there is a tradeoff with deployment simplicity. This approach would require a different `sudo` configuration for targets configured for different Delphix Engines.

Example: Configuring the `/etc/sudoers` File on the Target Environment for Privileges on the VDB Mount Directory Only (Linux OS)

```

Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount          <delphix-server-name>* /postgres/*, \
/bin/mount *        <delphix-server-name>* /postgres/*, \
/bin/mount          <delphix-server-ip>* /postgres/*, \
/bin/mount *        <delphix-server-ip>* /postgres/*, \
/bin/mount "", \
/bin/umount         /postgres/*, \
/bin/umount *       /postgres/*, \
/bin/mkdir [*]      /postgres/*, \
/bin/mkdir          /postgres/*, \
/bin/mkdir -p       /postgres/*, \
/bin/mkdir -p -m 755 /postgres/*, \
/bin/rmdir          /postgres/*

```

Related Topics

- [PostgreSQL Support and Requirements](#)

Managing PostgreSQL Environments and Hosts

This topic section the attributes of Postgres environments such as information for the Delphix Connector and covers the following topics:

- [Setting Up PostgreSQL Environments: An Overview](#)
- [Environment Attributes for Hosts with PostgreSQL](#)
- [Adding a PostgreSQL Environment](#)
- [Adding an Installation to a PostgreSQL Environment](#)
- [Changing the Hostname or IP Address for PostgreSQL Source and Target Environments](#)

Setting Up PostgreSQL Environments: An Overview


There are broadly three ways in which PostgreSQL source database can be made in sync with the staging Database:

1. Delphix Initiated Backup Ingestion + PostgreSQL Streaming Replication between Source and Staging Database.
2. Customer Initiated Backup Ingestion + PostgreSQL Streaming Replication between Source and Staging Database.
3. Customer Initiated Backup Ingestion + Source Database WAL logs application on Staging Server.


Each of the above approaches are described below:

Delphix for PostgreSQL Architecture with Delphix Initiated Backup Ingestion + PostgreSQL Streaming Replication between Source and dSource

Below is the brief architecture of Delphix Engine support for PostgreSQL database with Delphix Initiated Backup + PostgreSQL Streaming Replication between Source and Staging Database.




Block Diagram of Linking Architecture Between PostgreSQL Environments and the Delphix Initiated Backup Platform.




Delphix for PostgreSQL Architecture with Customer Initiated Backup Ingestion + PostgreSQL Streaming Replication between Source and dSource

Below is the brief architecture of Delphix Engine support for the PostgreSQL database with External Backup + PostgreSQL Streaming Replication between Source and Staging Database.



Delphix for PostgreSQL Architecture with Customer Initiated Backup Ingestion + WAL apply for resync between Source and dSource

Below is the brief architecture of Delphix Engine support for the PostgreSQL database with External Backup + WAL application between Source and Staging Database.



Types of PostgreSQL Environments

At a high level, the Delphix Engine maintains a logical representation of the source database files, from which one can provision virtual databases (VDBs). In order to link a data source and provision a VDB, the following types of environments are required:

A source environment is where the un-virtualized source database runs. The Delphix Engine uses the backup, restore, and replication features of the PostgreSQL DBMS to maintain its internal representation of the source database, to be used for provisioning VDBs. The Staging host must be able to connect to the source environment in order to orchestrate the backup, restore, and replication functionality necessary to keep its representation synchronized with the source database. The Delphix Engine is designed to have a minimal impact on the performance of the source database and the source environment.

A target environment is where virtualized databases run. PostgreSQL target environments serve two purposes:

1. Since PostgreSQL does not provide a native incremental backup API, a warm standby server (in other words, one in log-shipping mode) must be created where all database files are stored remotely on a Delphix Engine. We refer to the creation and maintenance of this staging database as validated sync. During validated sync, we retrieve data from the source, roll the staged database forward, ensure that all the components necessary for provisioning a VDB have been validated, and create a snapshot of that data state. The result of validated sync is both a TimeFlow with consistent points from which you can provision a VDB and a faster provisioning process because there is no need for any database recovery when provisioning a VDB. In order to create a staging database, you must designate a target environment for this task when linking a dSource. During the linking process, database files are exported over the network to the target environment, where

the staging database instance runs as a warm standby server. A target environment that hosts one or more staging databases is referred to as a staging target for validated sync.

2. Once a staging database has been set up, you can provision virtual databases from any of the discreet snapshots along the TimeFlow mentioned above to any compatible target environment (for more information, see [Requirements for PostgreSQL Hosts and Databases](#)). Database files are exported over the network to the target environment, where the virtual database instance runs.

Workflow for PostgreSQL Environments

Prior to provisioning a virtual database, a compatible target environment must be added to the Delphix Engine. This may be the same target environment as that used for the staging instance, or it may be a different target environment.

Once an environment is added to the Delphix Engine, environment discovery takes place. Environment discovery is the process of enumerating PostgreSQL installations and configurations when a source or target environment is added to the Delphix Engine. The discovery process is repeated during environment refresh in order to detect new PostgreSQL installations and clusters.

Related Topics

- [Managing PostgreSQL Environments and Hosts](#)

Environment Attributes for Hosts with PostgreSQL

This topic describes the attributes of PostgreSQL environments. Below you will see a section for common environment attributes shared by all types of environments as well as Postgres specific ones.

Procedure

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. In the Environments panel, click the name of an environment to view its attributes.
5. Next to **Attributes**, click the Pencil icon to edit an attribute.

Common Environment Attributes

Attribute	Description
Environment Users	The users for that environment. These are the users who have permission to ssh into an environment or access the environment through the Delphix Connector. For more information on the environment user requirements, see the Requirements topics for specific data platforms.
Host Address	The IP address of the environment host.
DSP KeyStore Path	The path to the user-managed DSP Keystore.
DSP KeyStore Alias	The lowercase alias to use inside the user-managed DSP Keystore.
DSP KeyStore Password	The password for the user-managed DSP Keystore.
DSP TrustStore Path	The path to the user managed DSP truststore.
DSP TrustStore Password	The password for the user managed DSP truststore.
OS	The name of the host operating system.
Version	The version of the host operating system.

Attribute	Description
Release	The release of the host operating system.
Time Zone	The timezone of the host operating system.
Total RAM	The amount of RAM on the host machine.
Processor Type	The processor type of the host machine.
Traceroute	Traceroute info from target host to Delphix Engine.
Notes	Any other information you want to add about the environment.
Java Development Kit	The currently selected JDK kit will be shown.
Java Development Kit (JDK) Path	Location of the Java Development Kit (JDK) used for the host. Only specified if the feature to provide your own JDK is enabled, otherwise, the defaults are used per our Java Support Matrix .

PostgreSQL Environment Attributes

There are no postgres-specific environment attributes.

Adding a PostgreSQL Environment

This topic describes how to add a PostgreSQL Staging environment.

Prerequisites

- Make sure that the staging environment in question meets the requirements described in [Requirements for PostgreSQL Hosts and Databases](#)

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next, to Environments, click the Actions (...) menu and select **Add Environment**.
5. In the Add Environment dialog, select **Unix/Linux**.
6. Select **Standalone Host**.
7. Click **Next**.
8. Enter Name for the Environment.
9. Enter the **Host IP address** or **hostname**.
10. Enter the **SSH port**. The default value is 22.
11. Enter an **OS Username** for the Environment which is currently 'postgres'.
12. Select **Login Type**.
 - Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

i Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

13. For Password, enter the password associated with the user in step 11.
14. If you want to use Public Key Encryption for logging into your environment:

- a. Select Public Key for the Login Type.
- b. Click View Public Key.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

The public key needs to be added only once per user and per environment.

You can also add public key authentication to an environment user's profile by using the command-line interface, as explained in the topic [CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users](#).

15. For Password Login, click Verify Credentials to test the username and password.
16. Enter **Toolkit Path** (make sure Toolkit path does not have spaces).
17. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
18. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
19. Click **Submit**.

As the new environment is added, you will see two jobs running in the Delphix platform Job History, one to Create and Discover an environment, and another to create an environment. When the jobs are complete, you will see the new environment added to the list in the Environments tab. If you do not see it, click the Refresh icon in your browser.

18. Once the environment is discovered, further linking would require adding a source config to the above-discovered installation. Please refer to [Linking a PostgreSQL dSource](#) for more information.

The screenshot displays the Delphix platform interface. On the left, a sidebar lists two environments: 'target' (Unix Host - Enabled) and 'target95' (Unix Host - Enabled). The main content area shows details for the 'target' environment, which is expanded to show 'Postgres vFiles (9.6)'. Under this section, there are two sub-sections: 'Installation Details' and 'Databases'. The 'Installation Details' section shows 'Allow Provisioning' as 'Yes', 'Repository Name' as 'Postgres vFiles (9.6)', 'Version' as '9.6', 'Postgres Install Path' as '/usr/pgsql-9.6', and 'Postgres Shell Path' as '/usr/pgsql-9.6/bin/psql'. The 'Databases' section shows 'No databases found on installation.' Below this, there is another section for 'Unstructured Files', which also shows 'Installation Details' with 'Allow Provisioning' as 'Yes' and 'name' as 'Unstructured Files'. The 'Databases' section for 'Unstructured Files' also shows 'No databases found on installation.'

Post Requirements

To view information about an environment after you have created it:

1. Click Manage.
2. Select Environments.
3. Select the environment name.

Related Topics

- [Managing PostgreSQL Environments and Hosts](#)

Adding an Installation to a PostgreSQL Environment

If a new PostgreSQL Installation is added to the environment, then in order to reflect it into the Delphix Engine, it is required to refresh the environment. Please refer to the "Environment Refresh" section at [Managing Environments and Hosts](#) for more details.

This topic describes how to refresh a PostgreSQL environment to reflect the changes if any.

Related Topics

- [Managing PostgreSQL Environments and Hosts](#)

Changing the Hostname or IP Address for PostgreSQL Source and Target Environments

Below is the procedure to change the Host Address of the PostgreSQL Source and Target Environments (If the Host address of the underlying Source and Target server has also changed):

Procedure

1. Disable the dSources.
2. Disable the VDBs.
3. Login to the **Delphix Management** application.
4. Click **Manage**.
5. Select **Environments**.
6. Click on the existing environment name you want to modify and open the environment information screen.
7. In the **Attributes** section, click the **Pencil** icon located next to it.
8. Update the **Host address** to the new one and save the results by clicking on the **Save (Tick)** icon.
The new Host address will now be updated.
9. Re-enable the dSources.
10. Re-enable the VDBs.

Related Topics

- [Setting Up PostgreSQL Environments: An Overview](#)

Linking Data Sources and Syncing Data with PostgreSQL

Creating a dSource will ingest data from the source and create a dSource on the engine. The dSource is an object that the Delphix Virtualization Engine uses to create and update virtual copies of your database. As a virtualized representation of your source data, it cannot be managed, manipulated, or examined by database tools.

This section covers the following topics:

- [Linking PostgreSQL Data Sources: Overview](#)
- [Data Management Settings for PostgreSQL Data Sources](#)
- [Linking a PostgreSQL dSource](#)
- [Working with PostgreSQL Snapshots](#)

Linking PostgreSQL Data Sources: Overview

This topic describes basic concepts behind the creation of a dSource from PostgreSQL instance.

For an overview of all dSource related actions, please visit [Managing Data Sources and Syncing Data](#) in the Getting Started section of the Datasets documentation.

When linking a dSource from a PostgreSQL source database, Delphix offers three different methods of capturing backup information:

1. dSource created with Delphix Initiated Backup, leveraging PostgreSQL Streaming Replication to sync data from the source database.
2. dSource created with External Backup, and syncing data with WAL files
3. dSource created with External Backup, and syncing data with PostgreSQL Streaming Replication

 If you are using an external backup and have checked both options above, then PostgreSQL Streaming Replication will take precedence, and syncing with WAL files will not be considered.

 For Delphix initiated and external backups, the streaming replication is monitored by reading the PostgreSQL database logs every minute, and if the streaming replication stops due to any reason, a fault is generated. An error message is displayed under the **Status tab**. The snapshot operation will also fail if the streaming replication stops with a fault.

After obtaining the initial snapshot and linking the dSource, the Delphix Engine keeps the staging and the source database in sync by monitoring the source database for new transaction logs on the staging target and then applying those transaction logs on the staging database. PostgreSQL streaming replication protocol is used to achieve data replication between the source database and the staging database.

 We can also use the source host as a staging host because the PostgreSQL database runs on different ports.

Related Topics

- [Linking Data Sources and Syncing Data with PostgreSQL](#)

Data Management Settings for PostgreSQL Data Sources

This topic describes advanced data management settings for dSources.

When linking a dSource, you can use custom data management settings to improve overall performance and match the needs of your specific server and data environment. If no specific settings are required, leverage default data management settings.

Accessing Data Management Settings

There are three ways to set or modify data management settings for dSources:

1. During the dSource linking process, in the Policies tab of the Add dSource wizard, select the **Retention Policy**.
2. Or select **Manage**, then select **Policies**. To manage retention on target after the dSource snapshots have been deleted from the replication source, select the **Replica Retention** tab.
3. In the Policies screen, select the **Retention** tab.
4. To add a new retention policy click the **+Retention** button, and to add a new replica retention policy on the replication target, click the **+Replica Retention** button.
5. This will open the Retention Policy wizard. Enter a policy name, select your retention periods and click Submit.
For more information, see [Policies for Scheduled Jobs](#).

Retention Policies - Retention/Replica Retention

Retention policies define the length of time Delphix Engine retains snapshots within its storage. To support longer retention times, you may need to allocate more storage to the Delphix Engine. The retention policy – in combination with the SnapSync policy – can significantly impact the performance and storage consumption of the Delphix Engine.

Replica Retention policies define how long the snapshots are retained on replicated namespaces for dSources and VDBs after they have been deleted on the replication source. Normally, the snapshots that have been deleted on the replication source engine are also deleted on the replication target engine. A new retention policy is introduced to provide an extended lifetime of such snapshots on the replication target.

Benefits of Longer Retention

With increased retention time for snapshots and logs, you allow a longer (older) rollback period for your data.

Common use cases for longer retention include:

- SOX compliance
- Frequent application changes and development
- Caution and controlled progression of data
- Reduction of project risk
- Restore to older snapshots

Related Topics

- [Linking Data Sources and Syncing Data with PostgreSQL](#)
- [Policies for Scheduled Jobs](#)

Linking a PostgreSQL dSource

Overview of the Linking Process

The PostgreSQL plugin version 1.1.0 and above provides the following ways to create a dSource:

- [dSource created with Delphix Initiated Backup](#) - this keeps in sync through PostgreSQL Streaming Replication
- dSource created with External Backup - there are two ways to do this:
 - keep in sync with the source through [WAL files](#)
 - keep in sync with the source through [PostgreSQL Streaming Replication](#)

 If you have checked both the options - **keep in sync with the source through WAL files** and **keep in sync with the source through PostgreSQL Streaming Replication**, then **keep in sync with the source through PostgreSQL Streaming Replication** will take precedence and syncing through WAL files will not be considered.

The following sections provide step by steps instructions on how to do this.

Creating a dSource with Delphix Initiated Backup

Prerequisites

- The source and staging instances must meet the host requirements
 - Databases must meet container requirements
- For more information refer to [Requirements for PostgreSQL Source Hosts and Databases](#)

Limitation

It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.

Procedure

1. Login to the **Delphix Management** application.
2. Select **Manage > Environments**.
3. From the **Databases** tab, select a repository from which you want to create your dSource and click the  icon.
4. In the **Add Database** dialog window enter the **Name** for your source config and click **Add**.
5. Select your source config and click the **Add dSource** link located to the right.
6. In the Source tab click on the **+Add** icon located next to **Delphix Initiated Backup** Parameter.

 The Add Button provided for Delphix Initiated Backup/External Backup - Streaming Replication is clickable multiple times but should be used **ONLY** once since the solution is built only to support a single backup path. Providing Multiple Paths will error out the linking process.

The following fields are mandatory for this process :

- **Delphix Initiated Backup Flag** - Checkbox should be selected.
- **PostgresDB Replication User**
- **PostgresDB Replication User Password**
- **Source Host Address**
- **Source Instance Port Number**
- **Staging Instance Port Number** - must be unique for a given instance

Data Type
AppData

External Backup

dSource Creation through external backup

[+ Add](#)

Allow the dSource to be in sync through streaming replication if WallLogs are not provided.
Allow the dSource to be in sync with Source through streaming replication if WallLogs are not provided. "Delphix Initiated Backup" parameters are mandatory for this.

Delphix Initiated Backup Flag
Allow Delphix to initiate the backups from source database. "Delphix Initiated Backup" parameters are mandatory for ingestion.

Delphix Initiated Backup/External Backup - Streaming Replication

List of parameters for Delphix initiated backup and external backup with streaming replication

PostgresDB Replication User *
delphix [Delete](#)

Replication User for postgres database

PostgresDB Replication User Password *

Replication password for postgres database

Source Host Address *
centos78pg12src.dlpxdc.co

Source Host Address

Source Instance Port Number *
5432

Port number for postgres source database

[+ Add](#)

Staging Instance Port Number *
5433

Port number for postgres staging database

Config Settings

Custom Database-Level config settings

[+ Add](#)

[Cancel](#) [Back](#) [Next](#) [Submit](#)

7. Optionally, database configuration parameters can be defined during the linking operation in the Config Settings section.

Config Settings

Custom Database-Level config settings

Delete

Comment Property

Select this option to comment out the provided property name in the configuration file

+ Add

i After dSource creation, users can now configure the parameters and their values in the `postgresql.conf` file through the "Config Settings" section on the Delphix Engine UI. For example, to support **Replication Slots** through UI, users can now provide the "primary_slot_name" parameter and its value through the UI. This will update the already existing parameter value in the `postgresql.conf` file with the value provided by the user.

Similarly, if a user wants to disable a parameter then they can provide the parameter name and select the check box **Comment Property**. This will comment out the parameter in the `postgresql.conf` configuration file.

Click **Next**.

8. In the **dSource Configuration** tab, enter a dSource name and select a group for your dSource. Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
9. In the **Data Management** tab specify your **Staging Environment**, **User**, and **Mount Base**, then click **Next**.
10. If required apply policy details to the dSource, then click **Next**.
11. Users can provide bash scripts as Plugin defined hooks. These hooks are executed after data ingestion but before snapshot operation, click **Next**.
12. Review the dSource Configuration and Data Management information in the **Summary** section.
13. Click **Submit**.

The Delphix Engine initiates two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking **Active Jobs** from the top menu bar, or by selecting **System > Event Viewer**. When the jobs have successfully completed, the database icon will change to a dSource icon on the **Environments > Host > Databases** screen, and the dSource will also appear in the list of Datasets under its assigned group.

⚠ The dSource Configuration Screen

After you have created a dSource, the dSource Configuration tab allows you to view information about it and make modifications to its policies and permissions. In the Datasets panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the Source files, Data

Management configuration, and Hook Operations. For more information, refer to [Data Management Settings for PostgreSQL Data Sources](#).

Creating a dSource with External Backups using WAL Logs

The following procedure provides details on creating a dSource with External Backup which uses WAL logs to keep in sync with the source.

Prerequisites

- The source and staging instances must meet the host requirements
- Databases must meet container requirements
For more information refer to [Requirements for PostgreSQL Source Hosts and Databases](#)
- PostgreSQL Backup Files Path - must be accessible by the OS user
- PostgreSQL Backup WAL Log Files Path - must be accessible by the OS user

Limitation

It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.

Procedure

1. Login to the **Delphix Management** application.
2. Select **Manage > Environments**.
3. From the **Databases** tab, select a repository from which you want to create your dSource and click the  icon.
4. In the **Add Database** dialog window enter the **Name** for your source config and click **Add**.
5. Select your source config and click the **Add dSource** link located to the right.
6. In the Source tab click on the **+Add** icon located next to **Delphix Initiated Backup** Parameter.

 The Add Button provided for External Backup is clickable multiple times but should be used **ONLY** once since the solution is built only to support a single backup path. Providing Multiple Paths will error out the linking process.

- **PostgreSQL Backup Files Path - the location of the backup file**

The command below is the pg_basebackup command to be used for taking backups of the PostgreSQL Database:

- **For PostgreSQL versions 9.4.x,9.5.x,9.6.x**
`<path to bin>/pg_basebackup -p<port number> -D<data dir location> -F t -v -w -P -x`
- **For PostgreSQL versions 10.x and 11.x**
`<path to bin>/pg_basebackup -p<port number> -D<data dir location> -F t -v -w -P -Xs`
- **PostgreSQL Backup WAL Log Files Path** - the location of the WAL log files

 Once the dSource is up using the Full External Backup provided on the staging environment, the dSource will continue to run in recovery mode and will wait for the next expected WAL file under the specified **PostgreSQL backup WAL log Files** path. When the expected WAL files are placed on this path, the dSource applies this WAL file and starts waiting for the next expected WAL files in series. If the WAL file provided is invalid or not in the sequence, then the dSource will shut down (INACTIVE) and will **ONLY** throw an error on UI with the information of expected WAL filename when next dSource snapshot is taken.

In the case of a WAL chain break, users need to either resynchronize the dSource or provide the missing WAL file, then disable and enable the dSource database

- **Staging Instance Port Number** - must be unique for a given instance.

Data Type
AppData

External Backup

dSource Creation through external backup

PostgreSQL Backup Files Path *
/tmp/backupfile Delete

PostgreSQL Backup Files Path

PostgreSQL Backup WAL Log Files Path
/tmp/wal_Logs

Allow the dSource to restore the WAL log files from this path

+ Add

Allow the dSource to be in sync through streaming replication if WalLogs are not provided.

Allow the dSource to be in sync with Source through streaming replication if WalLogs are not provided. "Delphix Initiated Backup" parameters are mandatory for this.

Delphix Initiated Backup Flag

Allow Delphix to initiate the backups from source database. "Delphix Initiated Backup" parameters are mandatory for ingestion.

Delphix Initiated Backup/External Backup - Streaming Replication

List of parameters for Delphix initiated backup and external backup with streaming replication

+ Add

Staging Instance Port Number *
5433

Port number for postgres staging database

Config Settings

Custom Database-Level config settings

+ Add

Cancel Back Next Submit

1. Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section. Click **Next**.
2. In the **dSource Configuration** tab, enter a dSource name and select a group for your dSource. Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its

objects, such as snapshots. See the topics under [Users and Groups](#) for more information.

Data Type
AppData

External Backup

dSource Creation through external backup

[+ Add](#)

Allow the dSource to be in sync through streaming replication if WalLogs are not provided.
Allow the dSource to be in sync with Source through streaming replication if WalLogs are not provided. "Delphix Initiated Backup" parameters are mandatory for this.

Delphix Initiated Backup Flag
Allow Delphix to initiate the backups from source database. "Delphix Initiated Backup" parameters are mandatory for ingestion.

Delphix Initiated Backup/External Backup - Streaming Replication

List of parameters for Delphix initiated backup and external backup with streaming replication

PostgresDB Replication User *
delphix [Delete](#)

Replication User for postgres database

PostgresDB Replication User Password *

Replication password for postgres database

Source Host Address *
10.110.11.12

Source Host Address

Source Instance Port Number *
5432

Port number for postgres source database

[+ Add](#)

Staging Instance Port Number *
5433

Port number for postgres staging database

Config Settings

Custom Database-Level config settings

[+ Add](#)

Cancel Back Next Submit

3. In the **Data Management** tab specify your **Staging Environment**, **User**, and **Mount Base**, then click **Next**.
4. If required apply policy details to the dSource, then click **Next**.
5. Users can provide bash scripts as Plugin defined hooks. These hooks are executed after data ingestion but before snapshot operation, click **Next**.
6. Review the dSource Configuration and Data Management information in the **Summary** section.
7. Click **Submit**.

⚠ The dSource Configuration Screen

After you have created a dSource, the dSource Configuration tab allows you to view information about it and make modifications to its policies and permissions. In the Datasets panel, select the

dSource you wish to examine. You can now choose the configuration tab to see information such as the Source files, Data Management configuration, and Hook Operations.

Creating a dSource with External Backups using Streaming Replication

The following procedure provides details on creating a dSource with External Backup which uses streaming replication to keep in sync with the source.

Prerequisites

- The source and staging instances must meet the host requirements
- Databases must meet container requirements
For more information refer to [Requirements for PostgreSQL Hosts and Databases](#)
- PostgreSQL Backup Files Path - must be accessible by the OS user

Limitation

It is not possible to access the staging server with PostgreSQL 9.4 and PostgreSQL 9.5 versions.

Procedure

1. Login to the **Delphix Management** application.
2. Select **Manage > Environments**.
3. From the **Databases** tab, select a repository from which you want to create your dSource and click the  icon.
4. In the **Add Database** dialog window enter the **Name** for your source config and click **Add**.
5. Select your source config and click the **Add dSource** link located to the right.
6. In the Source tab click on the **+Add** icon located next to **Delphix Initiated Backup** Parameter.

 The Add Button provided for External Backup is clickable multiple times but should be used **ONLY** once since the solution is built only to support a single backup path. Providing Multiple Paths will error out the linking process.

- **PostgreSQL Backup Files Path** - the location of the backup file
- **dSource Creation through external backup** - should be selected.
- Delphix Initiated Backup/External Backup - Streaming Replication parameters :
 - **PostgresDB Replication User**
 - **PostgresDB Replication User Password**
 - **Source Host Address**
 - **Source Instance Port Number**

- **Staging Instance Port Number** - must be unique for a given instance

Data Type
AppData

External Backup

dSource Creation through external backup

[+ Add](#)

Allow the dSource to be in sync through streaming replication if WalLogs are not provided.
Allow the dSource to be in sync with Source through streaming replication if WalLogs are not provided. "Delphix Initiated Backup" parameters are mandatory for this.

Delphix Initiated Backup Flag
Allow Delphix to initiate the backups from source database. "Delphix Initiated Backup" parameters are mandatory for ingestion.

Delphix Initiated Backup/External Backup - Streaming Replication

List of parameters for Delphix initiated backup and external backup with streaming replication

PostgresDB Replication User *
delphix [Delete](#)

Replication User for postgres database

PostgresDB Replication User Password *
.....

Replication password for postgres database

Source Host Address *
10.110.11.12

Source Host Address

Source Instance Port Number *
5432

Port number for postgres source database

[+ Add](#)

Staging Instance Port Number *
5433

Port number for postgres staging database

Config Settings

Custom Database-Level config settings

[+ Add](#)

[Cancel](#)
[Back](#)
[Next](#)
[Submit](#)

- Optionally, users can set the database configuration parameters during the linking operation in the **Config Settings** section. Click **Next**.
- In the **dSource Configuration** tab, enter a dSource name and select a group for your dSource. Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
- In the **Data Management** tab specify your **Staging Environment**, **User** and **Mount Base**, then click **Next**.
- If required apply policy details to the dSource, then click **Next**.
- Users can provide bash scripts as Plugin defined hooks. These hooks are executed after data ingestion but before snapshot operation, click **Next**.
- Review the dSource Configuration and Data Management information in the **Summary** section.
- Click **Submit**.

The Delphix Engine initiates two jobs to create the dSource: DB_Link and DB_Sync. You can monitor these jobs by clicking **Active Jobs** from the top menu bar, or by selecting **System > Event Viewer**. When the jobs have successfully completed, the database icon will change to a dSource icon on the **Environments > Host > Databases** screen, and the dSource will also appear in the list of Datasets under its assigned group.

 The dSource Configuration Screen

After you have created a dSource, the dSource Configuration tab allows you to view information about it and make modifications to its policies and permissions. In the Datasets panel, select the dSource you wish to examine. You can now choose the configuration tab to see information such as the Source files, Data Management configuration, and Hook Operations.

Related Topics

- [Linking Data Sources and Syncing Data with PostgreSQL](#)

Working with PostgreSQL Snapshots

Taking a snapshot creates a new snapshot entry in the dSource's Timeflow. The plugin-managed dSource allows you to select the parameters before taking the snapshot. These parameters are provided as an input to pre-snapshot and post snapshot functions. On the other hand, in the default snapshot, the parameters are pre-defined in the dSource environment.

Snapshot (Default)

This section lists the steps to take a snapshot and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog, select **Perform Snapshot**.
6. Click **View:** under **Timeflow** to verify the Snapshot you just created.
7. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
8. In the Delete Snapshot dialog, select **Delete**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Provisioning and Managing VDBs from PostgreSQL dSources

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment.

From a dSource, you can select a snapshot or point in time to create a VDB. PostgreSQL VDBs each have their own configuration settings as described in the following topics:

- [Provisioning PostgreSQL VDBs: Overview](#)
- [Upgrading PostgreSQL VDBs](#)
- [Provisioning a PostgreSQL VDB](#)
- [Database Permissions for Provisioned PostgreSQL VDBs](#)
- [Rewinding a PostgreSQL VDB](#)

Provisioning PostgreSQL VDBs: Overview

This topic describes the basic concepts involved with provisioning VDBs from PostgreSQL.

Virtual databases are a key data management concept for Delphix. In order to create or provision a virtual database, you will need a linked dSource from a source host and a compatible target environment, as described in the overview for [Managing Environments and Hosts](#) and [Overview of Requirements for PostgreSQL](#).

From a dSource, you can select a snapshot or point in time to create a VDB. PostgreSQL VDBs each have their own configuration settings as described in [Configuration Settings for PostgreSQL Virtual Databases](#)

For an overview of the high-level components involved in provisioning a PostgreSQL VDB refer to [Setting Up PostgreSQL Environments: An Overview](#).

Once you have provisioned a VDB, you can also take snapshots of it. As with the dSource snapshots, you can find these when you select the VDB in the Datasets panel. You can then provision additional VDBs from these VDB snapshots.

Dependencies

If there are dependencies on the snapshot, you will not be able to delete the snapshot free space; the dependencies rely on the data associated with the snapshot.

Related Topics

- [Setting Up PostgreSQL Environments: An Overview](#)
- [Provisioning and Managing VDBs from PostgreSQL dSources](#)

Upgrading PostgreSQL VDBs

This topic describes how to upgrade a PostgreSQL VDB to a higher version of PostgreSQL.

 Currently, only minor upgrades are supported. For example, from PostgreSQL version 12.2 to 12.7.

Procedure for VDB In-Place Upgrade

1. Remove any VDB Refresh Policy assigned to the VDB.
2. Upgrade the target PostgreSQL instance.
3. Refresh the target environment.

Related Topics

- [Provisioning and Managing VDBs from PostgreSQL dSources](#)

Provisioning a PostgreSQL VDB

You can take a new snapshot of the dSource by clicking the Camera icon on the dSource card. Once the snapshot is complete you can provision a new VDB from it. This topic describes how to provision a virtual database (VDB) from a POSTGRES dSource.

Prerequisites

- You will need to have linked a dSource from a staging instance, as described in [Linking a PostgreSQL dSource](#) or have created a VDB from which you want to provision another VDB
- You should have set up the POSTGRES target environment with necessary requirements as described in [PostgreSQL Support and Requirements](#)
- Make sure you have the required Instance Owner permissions on the target instance and environment
- The method for [Database Permissions for Provisioned PostgreSQL VDBs](#) is decided before the provisioning

Procedure

1. Navigate to **Manage**, and select **Datasets**.
2. Select a dSource and a snapshot from which you want to provision. Click the provision VDB icon to open the **provision VDB** wizard.
3. Select a target environment from the left pane, and an Installation to use from the dropdown list of available PostgreSQL instances on that environment.
4. Set the **Environment User** to be the Instance Owner. **Note:** The picking of instance owner is only possible if you have multiple environment users set on that host.
5. You will see the **Target Configuration** section where you need to specify **Mount Path**.

6. Optionally, set the database configuration parameters for the VDB. Users can disable a configuration parameter by selecting the check box Comment Property. This will comment out the parameter in the `postgresql.conf` configuration file. Click **Next**.
7. On the configuration page enter the PostgreSQL database name as the vFiles Name value.

8. Select a **Target Group** for the VDB and click the green **Plus** icon to add a new group, if necessary.
9. Select a **Snapshot Policy** for the VDB then click **Next**.
10. Specify any desired hook operations.
11. Review the Provisioning Configuration and Data Management information.
12. Click **Submit**.

Once the VDB provisioning has successfully completed, if the source and target instance ids are not the same, you may want to grant secadm and dbadm on the database to the target instance id. Please refer to [Database Permissions for Provisioned PostgreSQL VDBs](#) for more information.

Related Topics

- [Provisioning and Managing VDBs from PostgreSQL dSources](#)

Database Permissions for Provisioned PostgreSQL VDBs

This topic describes the database permissions on provisioned POSTGRES virtual instances

PostgreSQL Authentication

Authentication is the process by which the database server establishes the identity of the client, and by extension determines whether the client application (or the user who runs the client application) is permitted to connect with the database username that was requested.

PostgreSQL offers a number of different client authentication methods. The method used to authenticate a particular client connection can be selected on the basis of (client) host address, database, and user.

PostgreSQL database user names are logically separate from user names of the operating system in which the server runs. If all the users of a particular server also have accounts on the server's machine, it makes sense to assign database user names that match their operating system usernames. However, a server that accepts remote connections might have many database users who have no local operating system account, and in such cases, there need be no connection between database user names and OS user names.

Delphix postgresQL authentication

Delphix for PostgreSQL requires that the staging and target hosts must already have the necessary users and authentication systems created/installed on them. Delphix will neither create users nor change database passwords as part of the provisioning process. It is required to have 'trust' as a PostgreSQL authentication mechanism on the source with a target server.

The following section describes some important authentication methods used by PostgreSQL architecture.

Trust Authentication

When trust authentication is specified, PostgreSQL assumes that anyone who can connect to the server is authorized to access the database with whatever database user name they specify (even superuser names). Of course, restrictions made in the database and user columns still apply. This method should only be used when there is adequate operating-system-level protection on connections to the server.

Password Authentication

The password-based authentication methods are md5 and password. These methods operate similarly except for the way that the password is sent across the connection, namely MD5-hashed and clear-text respectively.

PostgreSQL database passwords are separate from operating system user passwords. The password for each database user is stored in the pg_authid system catalog. Passwords can be managed with the SQL commands CREATE USER and ALTER ROLE, e.g., CREATE USER foo WITH PASSWORD 'secret'. If no password has been set up for a user, the stored password is null and password authentication will always fail for that user.

There are some more authentication mechanism which is less often used, refer to the Postgresql link embedded for more insights on them.

- GSSAPI Authentication
- SSPI Authentication
- Ident Authentication
- Peer Authentication
- LDAP Authentication
- RADIUS Authentication

- [Certificate Authentication](#)
- [PAM Authentication](#)
- [BSD Authentication](#)

Related Topics

- [Provisioning and Managing VDBs from PostgreSQL dSources](#)

Rewinding a PostgreSQL VDB

Rewinding a VDB rolls it back to a previous point in its Timeflow and re-provisions the VDB. The VDB will no longer contain changes after the rewind point.

⚠ Although the VDB no longer contains changes after the rewind point, the rolled-over Snapshots and Timeflow still remain in Delphix and are accessible through the Command Line Interface (CLI). Refer to [CLI Cookbook: Rolling Forward a VDB](#) for instructions on how to use these snapshots to refresh a VDB to one of its later states after it has been rewound.

⚠ Delphix clones a new Timeflow from the closest Snapshot older than or equal to the rewind point. This creates a dependency between the new Timeflow and the parent Snapshot and Timeflow. The parent Snapshot and Timeflow cannot be deleted because of this dependency. The VDB must first be refreshed before the parent Snapshot and Timeflow can be removed.

Prerequisites

To rewind a VDB, you must have the following permissions:

A user with Delphix Admin credentials can perform a VDB rewind on any VDB in the system.

Procedure

1. Log in to the **Delphix Management** application.
2. Under **Datasets**, select the VDB you want to rewind.
3. On the **Timeflow** tab, select the rewind point as a snapshot or a point in time.
4. Click **Rewind**.
5. If you want to use login credentials on the target environment other than those associated with the environment user, click Provide Privileged Credentials.
6. Click **Yes** to confirm.

⚠ Using TimeFlow Bookmarks

You can use a Timeflow bookmark as the rewind point when using the CLI. Bookmarks can be useful to:

- Mark where to rewind to – for example, before starting a batch job on a VDB
- Provide a semantic point to revert back to in case the chosen rewind point turns out to be incorrect.

For a CLI example using a Timeflow bookmark, refer to [CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark](#).

Related Topics

- [Provisioning a PostgreSQL VDB](#)

PostgreSQL Hook Operation Notes

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at `/bin/sh`. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh "$ARG_ENVIRONMENT_VARIABLE"
"second argument in double quotes" 'third argument in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a `bash` binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command-line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
```

```
# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

Shell Operation Tips

Using `nohup`

You can use the `nohup` command and process backgrounding from resource in order to "detach" a process from the Delphix Engine. However, if you use `nohup` and process backgrounding, you MUST redirect `stdout` and `stderr`.

Unless you explicitly tell the shell to redirect `stdout` and `stderr` in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either `stdout` or `stderr`. Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your `RunCommand` operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The `RunExpect` operation executes an `Expect` script on a Unix environment. The `Expect` utility provides a scripting language that makes it easy to automate interactions with programs which normally can only be used interactively, such as `ssh`. The Delphix Engine includes a platform-independent implementation of a subset of the full `Expect` functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Management application and CLI to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an `ssh` session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

PostgreSQL Environment Variables

Operations that run user-provided scripts have access to environment variables. For operations associated with specific dSources or virtual databases (VDBs), the Delphix Engine will always set certain environment variables so that the user-provided script can use them to access the dSource or VDB.

dSource Environment Variables

Environment Variable	Description
<code>DLPX_DATA_DIRECTORY</code>	The mount path provided by the user on the UI
<code>USER</code>	The OS user used to link the dSource

VDB Environment Variables

Environment Variable	Description
<code>DLPX_DATA_DIRECTORY</code>	The mount path provided by the user on the UI
<code>USER</code>	The OS user used to provision the VDB

Related Topics

- [Hook Scripts for Automation and Customization](#)

Managing Logs Retention

Delphix PostgreSQL plugin provides a feature that helps users to manage the log rotation for both debug and error logs.

Delphix users can configure the logging mechanism by modifying the parameters defined in the *dlpx_pg_logrotate.conf* file that gets created or is already available within the log directory. This configuration file is located at `<DLPX_TOOLKIT_PATH>/postgres/logs/dlpx_pg_logrotate.conf`.

This configuration handles the following parameters:

Setting Maximum Size

Delphix users can set the parameter `MAX_FILE_SIZE_KB` available in the *dlpx_pg_logrotate.conf* file to set the maximum size of the active log file in KB. Once this limit is reached, the plugin will rotate the log. The default/maximum value of this parameter is 10240 KB/10 MB. This parameter accepts only a positive integer value, such as 10240.

Setting Number of Log Files to Retention

Delphix users can set a retention number for both debug and error logs using the parameters `NUM_OF_FILES_TO_KEEP_DEBUG` and `NUM_OF_FILES_TO_KEEP_ERROR`, respectively. The default value of these parameters is 50. That means, the maximum number of rotated log files that will be retained are 50 files and each file should be of a maximum size as defined in the above `MAX_FILE_SIZE_KB` parameter. If the number of files exceeds this value, then the old files will be deleted in a FIFO order.

dlpx_pg_logrotate.conf

```
MAX_FILE_SIZE_KB=10240
NUM_OF_FILES_TO_KEEP_DEBUG=50
NUM_OF_FILES_TO_KEEP_ERROR=50
```

HANA Environments and Data Sources

i Starting 6.0.8.0 and HANA plugin version 6.0.0, Delphix provides you with a vSDK-based solution. **Delphix highly recommends their users to use this HANA Staged (vSDK-based) plugin over the existing HANA Direct (Lua-based) plugin.** Currently, the Lua-based HANA plugin 4.3.1 and earlier versions are also supported. Please refer to [HANA Release Notes](#) for supported versions. For Lua-based HANA plugin documentation, refer to [Delphix Engine 6.0.7.0 Documentation](#) and the earlier versions.

This section covers the following topics:

- [SAP HANA Overview](#)
- [Delphix Architecture](#)
- [HANA 2.0 Plugin Installation](#)
- [HANA 2.0 Plugin Tools](#)
- [Quick Start Guide for HANA plugin](#)
- [HANA Support and Requirements](#)
- [Managing HANA Environments and Hosts](#)
- [Linking Data Sources and Syncing Data with HANA](#)
- [Provisioning and Managing VDBs from HANA dSources](#)
- [HANA Hook Operations](#)

SAP HANA Overview

SAP HANA is an in-memory database, columnar-oriented database management system developed by SAP. The key differentiators of HANA are its tight integration into the SAP application stack coupled with its columnar store, the latter of which is intended to enable customers to fold both OLTP and OLAP operations into a single database entity. In practical terms, this means an SAP HANA database stores both analytical and transactional data. Analysis of the data combination is possible on the fly, helping businesses make real-time analytical decisions.

SAP HANA as a database product continues to evolve dramatically. Key among recent changes has been the movement from a Single Database Container (SDC) to Multi-Tenant Database Containers (MDC), an architectural change that has become the exclusive architecture of HANA version 2.0.

Implications of the HANA column store

In a row-based relational database, data is stored by row, usually in large files. Data is then accessed from the database by reading and returning full rows of data.

Row

Country	Product	Sales
1	US Alpha	3000
2	US Beta	1250 <-----storing this line across----->
3	JP Alpha	400
4	UK Alpha	700

Column

Country	Product	Sales	
US	Alpha	3000	^
US	Beta	1250	
JP	Alpha	400	
UK	Alpha	700	v

storing each column up / down

Select

In a columnar database, the entire row does not need to be read, only the columns associated with the query. This columnar structure improves (read) efficiency and, therefore, query performance, particularly for extensive data sets.

Columnar data storage allows highly efficient compression because most of the columns contain only a few distinct values (compared to the number of rows). The data, therefore, has a smaller memory footprint and takes fewer processing cycles to interrogate.

Better parallel Processing: in a column store, data is already vertically partitioned. This means that operations on different columns can easily be processed in parallel. If multiple columns need to be searched or aggregated, each of these operations can be assigned to a different processor core.

Insert / Update

Storing data in columns instead of rows is challenging for workloads with many data modifying operations. Therefore, the concept of a differential buffer was introduced, where new entries are written to a differential buffer first. In SAP HANA, this is referred to as the delta store. In contrast to the main store, the delta store is optimized for inserts. At a later point in time and depending on thresholds, e.g. the frequency of changes and new entries, the data in the delta store is merged into the main store. This process is referred to as delta merge.

The delta merge process has important implications for filesystem-based table modifications. Delta merge on any table will reorganize all data in that table and as a corollary will re-write the entire table to different filesystem blocks. *The space savings Delphix can offer by tracking block changes between snapshots becomes far less effective or predictable when entire tables are reorganized as a function of delta merge.*

Related Topics

- [HANA Environments and Data Sources](#)

Delphix Architecture


The Delphix HANA plugin is based on a staging architecture, i.e. the linking process and SnapSync requires a staging environment. The staging environment is a HANA DB server host that leverages storage allocated directly on the Delphix Engine using NFS, and it must contain an instance of SAP HANA that is higher or matches the same version of the Source. To create a dSource on a staging server, a full backup would be required. The Recover command deposits its backup data onto the mounted storage area. The Delphix Engine snapshots the Staging server and the snapshots can be provisioned out to one or more target servers.

To create a dSource on a staging server, a full backup of one particular tenant database is required. With this solution, virtualization will be done at the tenant level only. This topic describes the high-level architecture of the Delphix HANA plugin solution.

This solution invokes HANA recovery of database backup files when a user selects to create a dSource on the staging server. The necessary backups/logs are created outside of Delphix and must exist before dSource creation.

Virtual databases can only be provisioned from a dSource (inheriting the state of the dSource).

The following is the high-level Delphix architecture with SAP HANA.



There are two mechanisms to create dSource on the staging server.

- Pull architecture, where the HANA plugin pulls the data into the Delphix provided mount location.
- Staging Push architecture, where the user pushes the data into the Delphix engine.

SAP HANA Plugin with Pull Architecture

In Delphix HANA with Pull architecture, dSource is created using External backups which must be a HANA Native backup (ie: no support for third-party Backup tools). With External backups, the first dSource snapshot must include a Full backup, but can optionally include one or more incremental or differential backups. Subsequent dSource snapshots can be based upon additional incremental or differential backups (This will require a full backup along with incremental/differential backups, as recovery in Hana always requires a full backup).

HANA plugin supports an External backup, also known as Native Backups. Native backup implies that the backups are written in native storage format to local storage without a third-party backup tool.

The dSource creation operation invokes the HANA database recovery of the relevant HANA backup files. HANA database recovery can be a lengthy process. The duration of recovery time is subject to the size of the database being recovered. After completion of HANA recovery, the HANA instance will contain the newly recovered dSource tenant.

VDB provision operations from a dSource snapshot do not require HANA database recovery and can occur very quickly.

Failure to follow this recommendation diminishes the key value proposition for HANA, which is agility. Below are the recommended tiers:

- **dSource:** A virtual HANA database constructed by recovery.
- **Virtual Database (VDB):** A copy of a specific dSource snapshot, cloned using Delphix technology. These copies are fast to create with a small initial storage footprint. VDBs are the recommended workbench for developers. Multiple VDBs can exist from a single dSource. Each VDB corresponds to a dSource snapshot.

SAP HANA Plugin with Staging Push Architecture

Starting Delphix Engine version 6.0.12.0 and HANA plugin version 6.2.0, a new data ingestion mechanism has been introduced that will help users to push data into the Delphix provided mount point on their own.


The previous versions of the HANA plugin were based only on the Pull approach and it was difficult to support the different ingestion mechanisms as it required a lot of effort to analyze, code, and develop the solution. With the Staging Push mechanism, data ingestion can be achieved with minimum effort.

Few points to be noted about the Staging Push mechanism.

- dSource creation creates an empty mount point on the staging host
- User managed backup and restore

Operational Best Practice

- Construct a dSource using the backup files. Construct multiple VDBs descended from the dSource.
- Create a meaningful object name.
- Create Dataset Groups that help delineate to which tier an object belongs.



- Segregate user targets (where VDBs reside) from the staging where HANA recovery must take place to create the dSource.
- Provision end-user VDBs, so that they can be created, refreshed, and rewound quickly.
- Layer VDBs to perform any common post-recovery operations to further streamline provisioning steps.
- Closely monitor and manage the quantity and lifespan of dSource and VDB snapshots, since each can consume additional Delphix storage.

Benefits of HANA Plugin

- Virtualization of individual tenant databases instead of the entire system.
- Ability to mix and match virtualized tenants across any HANA system.
- Reduced Storage footprint. VDBs can use significantly less storage when initially provisioned from dSource.
- Relaxes strict requirement of exact HANA version match between source and target, leveraging HANAs ability to recover an earlier version against a more recent version.
- Source backup files can be ingested from a surrogate host without impact to the true production source host.

Limitations of HANA Plugin

- The tenant's name should contain at least one alphanumeric character or one underscore. Also, it should not start with a digit.
- A single quote, double quote, and space characters are not allowed in the password.
- Datasets created using the Lua version cannot be upgraded to the vSDK version.
- Manual cleanup/unmounting would be required after **force delete** operation on any dataset.

HANA 2.0 Plugin Installation

The HANA 2.0 Plugin is provided as a zip file on the Delphix download site. Follow the steps below to download the plugin:

1. In the web browser, go to the Delphix [download](#) site.
2. Login to the download site using email and password credentials.
3. Navigate to the required version number of Delphix Engine.
4. Go to the Plugins > HANA > HANA 2.0 > Plugin_<version_number>.zip folder and download the zip file.

For initial and upgrade installations refer to the section [Delphix Engine Plugin Management](#) for further details.

Related Links

- [HANA Support/Compatibility Matrix](#)
- [Delphix Engine Plugin Management](#)

HANA 2.0 Plugin Tools

Delphix provides the **create_backup_metadata.sh** script, which is bundled with the plugin. This script is required to be executed by customer for customer-created HANA backups, which are to be used to create or refresh a dSource. This script can be found at:

```
<path>/scripts/create_backup_metadata.sh
```

This script must be run prior to creating or refreshing a dSource from a customer-created backup. The script is not required for Delphix-initiated backups. The script must be run on the source instance where the source tenant database exists. The script retrieves and persists important metadata about the specific backup set. This information will be used by the plugin during the dSource creation and will then be stored in the metadata section of the dSource.

Prerequisites

- The script assumes the user has the credentials of the tenant on the source HANA instance.
- The script must be copied to the source host as the operating system user should have access to hdbsql.
- The script must be executable.
- Correct backup files of the source must exist on the source host or the backup files should be accessible from the source host.
- The script should have the write permissions in the directory where the backups have been kept.

You can run the `create_backup_metadata.sh` script in the following two ways.

1. Interactive mode
2. Providing the complete list of parameters (Non-Interactive mode)

Interactive Mode

The mode allows you to input the values to the required/optional parameters one by one on the CLI. See the sample script run example below the Parameters table.

Parameters

The following table lists the optional and required parameters used by the script.

Parameter	Optional/Required	Description
Please enter the location of the Backup files:	Required	Location of the Backup files For example: <code>/usr/sap/HDB/HDB00/source_backup</code>

<p>If the above provided location of the Backup files is different than the original location where backups were taken then please enter original location, otherwise press enter(OPTIONAL):</p>	Optional	<p>This is the original location where the backup files were created during the dSource creation which can be used later for recovery purposes</p> <p>Provide this parameter when the location provided at the "<i>Please enter the location of the Backup files:</i>" parameter is different from the location where the backup files are kept now</p> <p>For example: <i>/usr/sap/HDB/HDB00/backup/data/DB_TEST1</i></p>
<p>Please enter the location of the Log files(OPTIONAL):</p>	Optional	<p>Location of the Log files</p> <p>You may use logs for the recovery</p> <p>For example: <i>/usr/sap/HDB/HDB00/source_log</i></p>
<p>If the above provided location of the log files is different than the original location where log files were created then please enter original location, otherwise press enter(OPTIONAL):</p>	Optional	<p>This is the original location where the log files were created during the dSource creation which can be used later for recovery purposes</p> <p>Provide this parameter when the location provided at the "<i>Please enter the location of the Backup files:</i>" parameter is different from the location where the log files are kept now</p> <p>For example: <i>/usr/sap/HDB/HDB00/backup/log/DB_TEST1</i> (Note that this is not a valid directory)</p>
<p>Please enter the Tenant Database Name:</p>	Required	<p>Name of the tenant database for which backups/logs were created</p> <p>For example: TEST1</p>
<p>Please enter the SID of the HANA database:</p>	Required	<p>SID of the tenant database for which backups/logs are present</p> <p>For example: HDB</p>
<p>Enter the Tenant User Name:</p>	Required	<p>Username to connect to the tenant database</p> <p>For example: SYSTEM</p>
<p>Tenant User Password:</p>	Required	<p>Password to connect to the tenant database</p>

Please enter the Instance Number:	Required	Instance number of the tenant database For example: 00
-----------------------------------	----------	---

Example

The following example shows an execution of the script with input parameters:

```

hdbadm@ip-10-110-244-207.delphix:/usr/sap/HDB/HDB00> ./create_backup_metadata.sh
Starting create_backup_metadata.sh
Please enter the location of the Backup files:/usr/sap/HDB/HDB00/source_backup

If the above provided location of the Backup files is different than the original location where backups were taken then please enter original location, otherwise press enter(OPTIONAL):/usr/sap/HDB/HDB00/backup/data/DB_TEST1

Please enter the location of the Log files(OPTIONAL):/usr/sap/HDB/HDB00/backup/log/DB_TEST1

If the above provided location of the log files is different than the original location where log files were created then please enter original location, otherwise press enter(OPTIONAL):

Please enter the Tenant Database Name:TEST1

Please enter the SID of the HANA database:HDB

Please enter the Tenant User Name:SYSTEM

Tenant User Password:

Please enter the Instance Number:00
    
```

Usage

```

./create_backup_metadata.sh
    
```

Providing the Complete List of Parameters

This mode allows you to input the parameter values in a specified order at the run time. See the sample script run example below the Parameters table.

Parameters

The following table lists the optional and required parameters in the order that in which you need to pass the parameter values along with the script.

Parameter	Optional/Required	Description
TENANT_DATABASE_NAME	Required	Tenant database name for which the script needs to be run
SID	Required	SID of the HANA database
TENANT_USER_NAME	Required	Username to connect to the tenant database
TENANT_USER_PASSWORD	Required	Password to connect to the tenant database
INSTANCE_NUMBER	Required	Instance number of the HANA database

DATA_BACKUP_PATH	Required	The location where the backups are present
LOG_BACKUP_PATH	Optional	The location where the logs are present
ORIG_DATA_BACKUP	Optional	Original location where these backups were created. Do not provide any value if backups are present at the same location where they were created
ORIG_LOG_BACKUP_PATH	Optional	Original location where logs were created. Do not provide any value if the logs are present at the same location where they were created

Example 1

The following example shows an execution of the script with input parameter values in the above-specified order.

```
hdbadm@target:/usr/sap/HDB/HDB00> ./create_backup_metadata.sh TEST HDB SYSTEM Delphix_123 00 /usr/sap/HDB/HDB00/source_backup
Starting create_backup_metadata.sh
Checking : /usr/sap/HDB/HDB00/source_backup

Running the script for TEST tenant
BACKUP ID of latest complete backup: 1633404668463
Metadata information of backup is stored in file: /usr/sap/HDB/HDB00/source_backup/backupTimeInfo
```



If all the mandatory parameters are not provided or parameters are not provided in the specified order, then the plugin will throw an error.

Example 2

The following example shows an error when all the mandatory parameter values are not provided.

```
hdbadm@target:/usr/sap/HDB/HDB00> ./create_backup_metadata.sh TEST HDB SYSTEM Delphix_123
Starting create_backup_metadata.sh

Incorrect number of parameters provided to the script. Please provide the below parameters to the script or do not pass any parameters which will invoke the script in interactive mode.

TENANT_DATABASE_NAME - Tenant Database name for which script needs to be run.
SID - SID of the HANA Database
TENANT_USER_NAME - User name to connect to the Tenant Database
TENANT_USER_PASSWORD - Password to connect to the Tenant Database
INSTANCE_NUMBER - Instance number of the HANA Database
DATA_BACKUP_PATH - Location where backups are present
LOG_BACKUP_PATH - Location where logs are present. This is an optional field.
ORIG_DATA_BACKUP - Original location where these backups were created. This is an optional field. Do not provide any value if Backups are present at the same location where they were created.
ORIG_LOG_BACKUP_PATH - Original location where logs were created. This is an optional field. Do not provide any value if Logs are present at same location
```

Related Topics

- [HANA Environments and Data Sources](#)

Quick Start Guide for HANA plugin

This quick start guide, which is excerpted from the larger User Guide, is intended to provide you with a quick overview of working with HANA database objects in the Delphix Engine. It does not cover any advanced configuration options or best practices, which can have a significant impact on performance. It assumes that you are working in a Lab/Dev setting and attempting to quickly test Delphix functionality. It assumes you will use the VMware Hypervisor.

Overview

In this guide, we will walk through deploying a Delphix Engine, starting with configuring Source and Target database environments on Windows servers. We will then create a dSource, and provision a VDB.

For purposes of the QuickStart, you can ignore any references to Replication or Masking.

Deploy OVA on VMware

Use the Delphix-supplied OVA file to install the Delphix Engine. The OVA file is configured with many of the minimum system requirements. The underlying storage for the install is assumed to be redundant SAN storage.

1. Download the OVA file from <https://download.delphix.com>. You will need a support login from your sales team or a welcome letter.
 - a. Navigate to the Delphix Product Releases/<Current Version>/Appliance Images page.
2. Login using the vSphere client to the vSphere server (or vCenter Server) where you want to install the Delphix Engine.
3. In the vSphere Client, click **File**.
4. Select **Deploy OVA Template**.
5. Browse to the OVA file.
6. Click **Next**.
7. Select a **hostname** for the Delphix Engine.
This hostname will also be used in configuring the Delphix Engine network.
8. Select the **data center** where the Delphix Engine will be located.
9. Select the **cluster** and the **ESX host**.
10. Select one (1) **data store** for the **Delphix OS**. This datastore can be **thin-provisioned** and must have enough free space to accommodate the 300GB comprising the Delphix operating system.
11. Select four (4) or more **data stores** for Database Storage for the Delphix Engine. The Delphix Engine will stripe all of the Database Storage across these VMDKs, so for optimal I/O performance, each VMDK must be equal in size and be configured **Thick Provisioned - Eager Zeroed**. Additionally, these VMDKs should be distributed as evenly as possible across all four SCSI I/O controllers.
12. Select the **virtual network** you want to use.
If using multiple physical NICs for link aggregation, you must use vSphere NIC teaming. Do not add multiple virtual NICs to the Delphix Engine itself. The Delphix Engine should use a single virtual network. For more information, see [Optimal Network Architecture for the Delphix Engine](#).
13. Click **Finish**.
The installation will begin and the Delphix Engine will be created in the location you specified.
14. Once the installation has completed, power on the Delphix Engine and proceed with the initial system configuration as described in [Setting Up Network Access to the Delphix Engine](#).

i If your source database is 4 TB, you probably need 4 TB of storage for the Delphix Engine. Add at least 4 data disks of similar size for the Delphix VM. For example: for a source database of 4 TB, create 4 VMDKs of 1 TB each.

i For a full list of requirements and best practice recommendations, see [Virtual Machine Requirements for VMware Platform](#).

Setup Network Access to the Delphix Engine

1. Power on the Delphix Engine and open the Console.
2. Wait for the Delphix Management Service and Delphix Boot Service to come online. This might take up to 10 minutes during the first boot. Wait for the large orange box to turn green.
3. Press any key to access the sysadmin console.
4. Enter `sysadmin@SYSTEM` for the username and `sysadmin` for the password.
5. You will be presented with a description of available network settings and instructions for editing.

Delphix Engine Network Setup

To access the system setup through the browser, the system must first be configured **for** networking in your environment. From here, you can configure the primary **interface**, DNS, hostname, and **default** route. When DHCP is configured, all other properties are derived from DHCP settings.

To see the current settings, run `"get."` To change a property, run `"set <property>=<value>."` To commit your changes, run `"commit."` To exit **this** setup and **return** to the standard CLI, run `"discard."`

```

    defaultRoute    IP address of the gateway for the default route -- for
                    example, "1.2.3.4."

    dhcp            Boolean value indicating whether DHCP should be used for
                    the primary interface. Setting this value
                    to "true" will cause all other properties (address,
                    hostname, and DNS) to be derived from the DHCP
                    response

    dnsDomain       DNS Domain -- for example, "delphix.com"

    dnsServers      DNS server(s) as a list of IP addresses -- for example,
                    "1.2.3.4,5.6.7.8."

    hostname        Canonical system hostname, used in alert and other logs --
                    for example, "myserver"

    primaryAddress  Static address for the primary interface in CIDR notation
                    -- for example, "1.2.3.4/22"

```

Current settings:

```

defaultRoute: 192.168.1.1
dhcp: false

```

```
dnsDomain: example.com
dnsServers: 192.168.1.1
hostname: Delphix
primaryAddress: 192.168.1.100/24
```

6. Configure the `hostname` . If you are using DHCP, you can skip this step.

```
delphix network setup update *> set hostname=<hostname>
```

 Use the same `hostname` you entered during the server installation.

7. Configure DNS. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set dnsDomain=<domain>
delphix network setup update *> set dnsServers=<server1-ip>[,<server2-ip>,...]
```

8. Configure either a static or DHCP address.

DHCP Configuration

```
delphix network setup update *> set dhcp=true
```

Static Configuration

```
delphix network setup update *> set dhcp=false
delphix network setup update *> set primaryAddress=<address>/<prefix-len>
```

 The static IP address must be specified in CIDR notation (for example, `192.168.1.2/24`)

9. Configure a default gateway. If you are using DHCP, you can skip this step.

```
delphix network setup update *> set defaultRoute=<gateway-ip>
```

10. Commit your changes. Note that you can use the `get` command prior to committing to verify your desired configuration.

```
delphix network setup update *> commit
Successfully committed network settings. Further setup can be done through the browser at:
```

```
http://<address>
```

Type "exit" to disconnect, or any other commands to **continue** using the CLI.

11. Check that you can now access the Delphix Engine through a Web browser by navigating to the displayed IP address, or hostname if using DNS.
12. Exit setup.

```
delphix> exit
```

Setting Up the Delphix Engine

Once you set up the network access for Delphix Engine, navigate to the Delphix Engine URL in your browser for server setup.

The welcome screen below will appear for you to begin your Delphix Engine setup.

DELPHIX SETUP
Setup Help

Virtualization Setup

- Welcome
- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization
- Registration
- Summary

Welcome

Choose engine type to setup:

Virtualization
 Masking

This wizard will step you through the setup. During this process you will complete the following:

- Create your password for the default "sysadmin" user
- Set the system time
- Configure network and services
- Configure the storage pool
- Configure proxies, SMTP, and LDAP (these are optional)
- Register your software

After setup is complete, you will have two administrators defined:

- The system administrator, "**sysadmin**" with the password you defined. This will be the system administrator for the instance.
- The engine administrator, "**admin**" with the password you defined. This is typically a DBA who will administer all the data managed by the instance.

When setup is complete, log in as engine administrator to begin using your engine.

Back Next Submit

The setup procedure uses a wizard process to take you through a set of configuration screens:

- Administrators
- Time
- Network
- Network Security
- Storage
- Outbound Connectivity
- Authentication
- Network Authorization

- Registration
- Summary

1. Connect to the Delphix Engine at `http://<Delphix Engine>/login/index.html#serverSetup`.

The **Delphix Setup** application will launch when you connect to the server.

Enter your **sysadmin** login credentials, which initially defaults to the username **sysadmin**, with the initial default password of **sysadmin**. On first login, you will be prompted to change the initial default password.

2. Click **Next**.

Administrators

The Delphix Engine supports two types of administrators:

- System Administrator (**sysadmin**) - this is the engine system administrator. The sysadmin password is defined here.
- Engine Administrator (**admin**) - this is typically a DBA who will administer all the data managed by the engine.

On the Administrators tab, you set up the sysadmin password by entering an email address and password. The details for the admin are displayed for reference.

 The default domain user created on Delphix Engines from 5.3.1 is known as **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then disabling the `delphix_admin`.

Time

The engine time is used as the baseline for setting policies that coordinate between virtual databases and external applications

Choose your option to setup system time in this section. For a Quick Start, simply set the time and your timezone. You can change this later.

Network

The initial out-of-the-box network configuration in the OVA file is set to use a single VMXNET3 network adapter.

You have already configured this in the initial configuration. Delphix supports more advanced configurations, but you can enable those later.

Storage

You should see the data storage VMDKs or RDMs you created during the OVA installation. Click **Next** to configure these for data storage.

Serviceability

Choose your options to configure serviceability settings.

For a Quick Start, accept the defaults. You can change this later.

Authentication Service

Choose your options to configure authentication services.

For a Quick Start, accept the defaults. You can change this later.

Registration

If the Delphix Engine has access to the external Internet (either directly or through a web proxy), then you can auto-register the Delphix Engine:

1. Enter your **Support Username** and **Support Password**.
2. Click **Register**.

If external connectivity is not immediately available, you must perform manual registration.

1. Copy the **Delphix Engine registration code** in one of two ways:
 - a. Manually highlight the registration code and copy it to clipboard. Or,
 - b. Click **Copy Registration Code to Clipboard**.
2. Transfer the Delphix Engine's registration code to a workstation with access to the external network Internet. For example, you could e-mail the registration code to an externally accessible e-mail account.
3. On a machine with access to the external Internet, please use your browser to navigate to the Delphix Registration Portal at <http://register.delphix.com>.
4. Login with your Delphix support credentials (username and password).
5. Paste the **Registration Code**.
6. Click **Register**.

 Although your Delphix Engine will work without registration, we strongly recommend that you register each engine as part of the setup. Failing to register the Delphix Engine will impact its supportability and security in future versions.

To regenerate the registration code for a Delphix Engine please refer to, [Regenerating the Delphix Engine Registration Code](#). Delphix strongly recommends that you regenerate this code and re-register the engine regularly to maximize the Support Security of the Delphix Engine. Delphix recommends doing this every six months.

Summary

The final summary tab will enable you to review your configurations for System Time, Network, Storage, Serviceability, and Authentication.

1. Click the **Back** button to go back and to change the configuration for any of these server settings.
2. If you are ready to proceed, then click **Submit**.
3. Click **Yes** to confirm that you want to save the configuration.
4. Click **Setup** to acknowledge the successful configuration.
5. There will be a wait of several minutes as the Delphix Engine completes the configuration.

Requirements for HANA Staging/Target Host and Databases

Backup of the tenant must be present which needs to be recovered. Along with backup, metadata information generated by script `create_backup_metadata.sh` is also expected in the backup location. As we will be recovering a tenant through backup files the important thing while creating the dSource which should be taken care of is the user should have access to the backup directory where the backup file resides.

General staging Host Requirements

- HANA database must be installed and running.
- Read-write permission on the backup directory.
- There must be an operating system user which has HANA administrative user privileges `<sid>adm`:
- The Delphix Engine must be able to make an SSH connection to the staging environment using the operating system user.

- The operating system user must have read and execute privileges on the HANA binaries installed on the staging environment.
- The operating system user must have read, write, and execute access to the HANA data directories on the staging environment.
- A plugin directory must exist on the staging host. The directory must have the following properties:
 - Should be writable by the operating system user described above.
 - Should have at least 256 MB of available storage.
- TCP/IP connectivity to and from the staging environment must be configured as described in [General Network and Connectivity Requirements](#).
- Hostname and IP address must be correctly configured.

Requirements for HANA Target Host and Databases

A **target environment** is where virtualized databases run. A target environment must have a HANA instance installed and configured. The VDB will become a tenant on the target database instance.

General Target Environment Requirements

1. The operating system of the target environment must match the staging environment.
2. The SAP HANA installation on the target environment must have the same or higher version than the staging environment.
3. There must be an operating system user with the following privileges:
 1. The Delphix Engine must be able to make an SSH connection to the target environment using the operating system user.
 2. The operating system user must have read and execute privileges on the HANA binaries installed on the target environment.
 3. The operating system user must have permission to run mount and unmount as the superuser via sudo with neither a password nor a TTY. See [Sudo Requirements for the HANA 2.0 Plugin](#).
 - Prior to the discovery, a plugin directory on the target environment must exist with the following properties:
 1. The plugin directory must be writable by the operating system user mentioned above.
 2. The plugin directory must have at least 256 MB of available storage.
 - There must be a mount point directory (for example, /mnt) that will be used as the base for mount points that are created when provisioning a VDB with the following properties:
 1. The mount point directory must be writable by the operating system user mentioned above.
 2. The mount point directory should be empty.
 - TCP/IP connectivity to and from the source environment must be configured as described in [General Network and Connectivity Requirements](#).
 - The hostname and IP address must be correctly configured.
 - HANA instance on which the VDB creation is being done should have the empty ports to host the services for the tenant database. The default port number range for tenant databases is 3<instance>40–3<instance>99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances. You do this by configuring the property [multidb] reserved_instance_numbers in the global.ini file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040–30199 if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040–30299 if the first instance is 00). And so on.

Adding a HANA Environment

Prerequisites

- Make sure that the HANA environment in question meets the requirements described in [Requirements for HANA Staging Hosts and Databases](#) and [Requirements for HANA Target Hosts and Databases](#).

Procedure

- Login to the **Delphix Management** application.
- Click **Manage**.
- Select **Environments**.
- Next to **Environments**, click the **Actions (...)** menu and select **Add Environment**.
- In the Add Environment dialog, In the **Host OS** select **Unix/Linux** on the menu.
- In **Server Type**, select **Standalone Host**.
- Click **Next**.
- Enter **Name** for the environment.
- Enter the **Host IP address** or **fully qualified hostname**.
- Enter the **SSH port**.
The default value is 22.
- Enter an **OS Username** for the environment which should be '<sidadm>'.
Select **Login Type**.

 Currently, the Enterprise Password Vault option is not supported.

- Username and Password - enter the OS username and password
- Username and Public Key - enter the OS username.
- Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

1. Select **Public Key** for the **Login Type**.
2. Click **View Public Key**.
3. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - a. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - b. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - c. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

1. For Password, enter the password associated with the user in step 11.
2. Select **Public Key** for the **Login Type**.
3. Click **View Public Key**.

4. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - a. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - b. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - c. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.
5. For Password Login, click 'Validate' to test the username and password.
6. Enter a **Plugin Path**. The location specified must be a string without spaces. Ensure this directory path exists on the host in question. Environment discovery will fail if the directory does not exist on the host.
7. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
8. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**. For more information, see [Providing Your Own Oracle Java](#).
9. Click **Submit**. On success, the new environment will be added to the list in the Environments tab.
10. Select the "Databases" tab to view discovered HANA databases:

The screenshot shows a web interface with two tabs: 'Details' and 'Databases'. The 'Databases' tab is active. Below the tabs is a section titled 'Installations' with a filter dropdown set to 'none' and an 'Add Dataset Home' button. A dropdown menu is open, showing a selected installation: 'HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00'. Below this, there are two sub-sections: 'INSTALLATION DETAILS' and 'DATABASIS'. The 'INSTALLATION DETAILS' section lists various configuration parameters, while the 'DATABASIS' section shows a message: 'No databases found on installation.'

INSTALLATION DETAILS	DATABASIS
<p>Allow Provisioning Yes</p> <p>System Administrator User hdbadm</p> <p>SAP Mount Path /hana/shared</p> <p>HANA Version 2.00.052.00.1599235305</p> <p>Hostname source</p> <p>HANA Instance Number 00</p> <p>HANA Installation Path /usr/sap/HDB/HDB00</p> <p>HANA SID HDB</p> <p>Repository Name HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00</p>	<p>No databases found on installation.</p>

The screenshot shows the configuration for a HANA data source in the Delphix Management application. The interface is divided into two tabs: 'Details' and 'Databases'. The 'Databases' tab is active, showing a list of installations. The selected installation is 'HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00'. The configuration is split into two columns: 'INSTALLATION DETAILS' and 'DATABASES'. The 'INSTALLATION DETAILS' column includes fields for 'Allow Provisioning' (Yes), 'System Administrator User' (hdbadm), 'SAP Mount Path' (/hana/shared), 'HANA Version' (2.00.052.00.1599235305), 'Hostname' (source), 'HANA Instance Number' (00), 'HANA Installation Path' (/usr/sap/HDB/HDB00), 'HANA SID' (HDB), and 'Repository Name' (HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00). The 'DATABASES' column includes 'Allow Linking' (Enabled) and 'Source Config Name (Required)' (HDB). There are edit icons for several fields and a '+ Add dSource' button.

Linking a HANA Data Source

This topic describes basic concepts behind the creation of a dSource for a HANA database using conventional pull and Staging push mechanisms.

⚠ For Staging Push, only native/third-party backups are used to populate data in the staging database.

For Third-party Backups, users need to provide a “Backup Path” containing FULL and Incremental backups. These files will be copied to the Delphix Engine during the dSource linking process. In this scenario, the Delphix Engine will not require any credentials of the source database. Although credentials for the tenant database (for which backup has been placed) and systemdb of staging instances would be required.


Prerequisites for Customer provided Backups Ingestion Mechanism

- Users need to run the `create_backup_metadata.sh`. This script will create the `backupTimeInfo` file at the same location where backups are kept. This script will gather all the required metadata information and will store that information in the `backupTimeInfo` file. This information will then get used by the plugin at the time of creating dSource.
- More details regarding the script have been included in the [HANA 2.0 Plugin Tools](#).


Procedure

1. Login to the Delphix Management application.

2. Select **Manage > Environments**. Then select the Datasets tab.




3. Create a source configuration on which to build the dSource. You can do this by clicking on the + symbol. The Source configuration can have any name.
4. In the **Add dSource** wizard, select the required source configuration.
5. After selecting the source, below are the parameters required to create the dSource.
 - a. Parameters required for the conventional **pull** mechanism:
 - i. Mount Location on Staging Host
 - ii. SAP HANA Source Tenant Database Name
 - iii. Source SystemDB User Name
 - iv. Source SystemDB Password
 - v. Source Tenant Database User Name
 - vi. Source Tenant Database Password
 - vii. Customer Provided Data Backup Location: This represents the path to the location of the customer-provided backup files.
 - viii. Customer Provided Log Backup Location: This represents the path to log the location of the customer-provided log backup files. This parameter is optional.



- b. Parameters required for the **Staging push** mechanism
 - i. Mount Location on Staging Host
 - ii. SAP HANA Source Tenant Database Name
 - iii. Enable the **Staging Push** feature by selecting the checkbox "**Use Staging Push**". With this feature, the user handles the Source Tenant database and its state on their own and you don't


need to specify any other staging database parameters, data backup, and log backup locations.

If this check box is selected, the HANA plugin will not invoke the tenant creation and recovery operation and if this checkbox is not selected then the default pull approach will be used.



c. Click **Next**.

6. Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
7. Configure the SnapSync and Retention Policy. The SnapSync policy can be set to 'None' for customers who wish not to run a scheduled SnapSync job.
8. Optionally, users can provide the bash script as a plugin-defined hook. Both Pre-Snapshot and Post-SnapShot hooks can be used.
9. Select **Next**.
10. At the end of the workflow, review the dSource Configuration in the Summary pane as shown below:



Once the dSource linking job is submitted it is possible to monitor progress by selecting **Active Jobs** in the top menu bar. On successful completion, the new dSource will be visible in the list of Datasets under its assigned group.

After dSource creation using the Staging Push mechanism, an empty snapshot will be created but users can not create VDB's from the same. Also, before creating the subsequent snapshots, users will have to create the database and point the database to use the Delphix provided mount point.

Provisioning a HANA VDB

Prerequisites

- The provision requires a linked dSource built off a physical source tenant, as described in [Linking a HANA dSource](#), or a VDB.
- The HANA target environment should adhere to the [Requirements for HANA Hosts and Databases](#).
- Operating system users and database users should be configured.

VDB creation UI

The screenshot displays the 'Target Configuration' step of the VDB creation process. The sidebar on the left shows the progress through various steps, with 'Target Configuration' currently active. The main configuration area includes several required fields: 'NFS Mount Location', 'SAP HANA Target Tenant Database Name', 'Target SystemDB User Name', 'Target SystemDB Password', 'Target Tenant Database ADMIN User Name', and 'Target Tenant Database ADMIN Password'. Each field has a descriptive tooltip. The 'Configure Tenants Service and Port' section allows for adding multiple tenant services, with two examples shown: 'indexserver:30040' and 'xsengine:30043'. Each entry has a 'Delete' button, and there is an '+ Add' button to add more. At the bottom right, there are 'Cancel', 'Back', 'Next', and 'Submit' buttons.

Procedure

1. Log in to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource**.
5. Select a **snapshot** from which you want to provision.

i If the first or empty dSource Snapshot that was created using the Staging Push mechanism is detected by the plugin and an error message is displayed. This check is done for every VDB creation operation.

6. Click the **Provision VDB** icon to open the Provision VDB wizard.
7. Select a target environment from the left pane.
8. The HANA 2.0 Plugin supports one installation on the same host, hence the value in the **Installation** dropdown will be already populated. Multiple drop-down options should not be available.
9. Set the **Environment User** to be the Instance Owner. Note: The picking of instance owner is only possible if you have multiple environment users set on that host.
10. Click **Next**.
11. In the **Target Configuration** pane you will need to specify the following:

 ****** indicates the mandatory fields.

- a. **NFS Mount Location ***: Delphix recommends setting the trailing directory to match the VDB name to delineate mount points and support administration and troubleshooting. In the example shown, the **NFS Mount Location** is /mnt/provision/dev where dev matches the **SAP HANA Target Tenant Database Name**.
- b. **SAP HANA Target Tenant Database Name ***: This refers to the target tenant database name which will be created upon VDB creation.
- c. **Target SystemDB User Name ***: This refers to the SystemDB username in the target database.
- d. **Target SystemDB Password ***: This refers to the SystemDB password in the target database.
- e. **Target Tenant Database ADMIN User Name ***: This refers to the username of the tenant system user.
- f. **Target Tenant Database ADMIN Password ***: This refers to the password of the tenant system user.

 The SYSTEM username and password derived from the physical source database should be entered during the provision of a VDB because those credentials will be recovered during provision.

- g. **Configure Tenants Service and Port**: This refers to the list of tenant services and port numbers in the following format.

```
<service>:<port>
For example, indexserver:30043
```

You can click "**+Add**" to add more services.

-  • This is not a mandatory field. If a user doesn't provide a value to this, then the plugin will proceed with the above fields and it will allocate the service and port on its own.
- If an invalid service is provided, the plugin will proceed with input, but will eventually fail while creating the service and will display an error on the Delphix Engine UI.

12. In the configuration section
 - a. Set the vFile Name or accept the default, which is a random string. Delphix strongly recommends applying a consistent naming convention. We suggest this name should match the Target Tenant Name from the previous pane. Maintaining a consistent naming convention will support administration and troubleshooting activities. In the example shown, the vFile Name “dev” matches the actual tenant name and mount point from the previous pane.
 - b. Select a **Target Group** for the VDB.
Click the **Add Dataset Group** icon to add a new group if required.
 - c. Set **Auto VDB Restart** option to automatically restart the VDB, where the target host is rebooted. Delphix recommends enabling this setting.
13. Click **Next**.

14. **VDB Snapshot Policy.** Select the Snapshot policy if an automated VDB snapshot schedule is required. Optionally, the Snapshot policy may be disabled by selecting “None”.
15. Click **Next**.
16. **Masking.** This page offers the opportunity to link the provision job to a masking job. At this time, Delphix Masking does not support HANA. Please do not use this option. “Mask this VDB” should remain deselected.
17. **Hooks.** Run a customized shell script via a hook operation at a particular moment in the provision operation.
18. Click **Next**.
19. **Summary.** Presents an opportunity to review all the information as a part of the initial provision workflow.
20. Click **Submit**.
Once a provision job is submitted it is possible to monitor its progress by selecting “Active Jobs” in the top menu bar. On successful completion, VDB will be visible in the list of Datasets under its assigned group.

Next Steps

Congratulations! You have provisioned your first virtual database!

Now, perform some simple functional tests with your application. You can connect your app to the VDB using standard TNS/JDBC techniques. Delphix has already registered the VDB for you on the target listener.

We suggest the following next steps:

1. Drop a table and use the VDB Rewind feature to test the recovery of your VDB.
2. Take a snapshot of your dSource and refresh your VDB to quickly get fresh production data.
3. Provision a new VDB from your VDB to test the sharing data quickly with other teams.
4. Mask your new VDB to protect sensitive data. Provision of new VDBs from that masked VDB to quickly provide safe data to development and QA teams.

Related Topics

- [Linking Data Sources and Syncing Data with HANA](#)
- [Provisioning and Managing VDBs from HANA dSources](#)

HANA Support and Requirements

This section covers the following topics:

- [HANA Support/Compatibility Matrix](#)
- [Requirements for HANA Staging Hosts and Databases](#)
- [Requirements for HANA Target Hosts and Databases](#)
- [Network and Connectivity Requirements for HANA Environments](#)
- [Sudo Requirements for the HANA 2.0 Plugin](#)

HANA Support/Compatibility Matrix

This section describes the HANA (DBMS) versions that are supported by the HANA 2.0 Plugin, as well as the compatible operating systems (OS), for use on the source and target environments.

 Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.

Supported DBMS Versions

HANA Version (Standard and Enterprise Editions)	Support Package Stack (SPS)
2.0	04
2.0	05

 **Important:** A target may have a higher SPS release than a source; the reverse is not supported.

Key:

Color	Supported?
	Yes
	No

Supported Operating System

Supported OS Version	Supported DBMS Version	
	HANA 2.0 SPS 04 (Standard and Enterprise Editions)	HANA 2.0 SPS 05 (Standard and Enterprise Editions)
RHEL 8.1		
RHEL 7.3/7.4/7.5		
CentOS 7.3/7.4/7.5		
RHEL 7.6		

CentOS 7.6	✓	✓
SUSE 12 SP 03	✓	✗
SUSE 15 SP 02	✗	✓

Plugin/Delphix Engine Compatibility

The plugin should be installed on compatible Delphix Engines per the table below:

Delphix Engine	Lua-based plugin						vSDK-based Plugin						
	4.1.1	4.1.2	4.2.0	4.3.0	4.3.1	4.3.2	6.0.0	6.0.1	6.1.0	6.1.1	6.2.0	6.3.0	
6.0.0.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.1.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.2.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.3.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.4.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.5.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.6.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.7.x	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	
6.0.8.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.9.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.10.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	
6.0.11.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	

	Lua-based plugin						vSDK-based Plugin					
6.0.12.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
6.0.13.x	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Unsupported Features

- Multiple duplicate services are not supported in this release.
- Point-in-Time recovery of the HANA database is not possible.
- HANA 1.0 MDC and SDC configurations are not supported.
- Scale-out source and target environments are not supported.
- High isolation mode is not supported with this release.
- Partitioning is not supported with this release.
- V2P is not supported with this release.
- Although the Staging Push solution is generic and should work with HANA Streaming Replication (HSR) but for this release, we have not validated the same.
- Third-party backups tools are not supported with Pull based ingestion mechanism.
- More than one instance of the same HANA service is not supported with this release.
- Switching between the Staging Push to Pull model is not supported with this release.

Related Topics

- [HANA Support and Requirements](#)

Requirements for HANA Staging Hosts and Databases

Backup of the tenant must be present which needs to be recovered. Along with backup, metadata information generated by script `create_backup_metadata.sh` is also expected in the backup location. As we will be recovering a tenant through backup files the important thing while creating the dSource which should be taken care of is the user should have access to the backup directory where the backup file resides.

General staging Host Requirements

- HANA database must be installed and running.
- Read-write permission on the backup directory.
- There must be an operating system user which has HANA administrative user privileges `<sid>adm`:
- The Delphix Engine must be able to make an SSH connection to the staging environment using the operating system user.
- The operating system user must have read and execute privileges on the HANA binaries installed on the staging environment.
- The operating system user must have read, write, and execute access to the HANA data directories on the staging environment.
- A plugin directory must exist on the staging host. The directory must have the following properties:
 - Should be writable by the operating system user described above.
 - Should have at least 256 MB of available storage.
- TCP/IP connectivity to and from the staging environment must be configured as described in [General Network and Connectivity Requirements](#).
- Hostname and IP address must be correctly configured.

Related Topics

- [HANA Support and Requirements](#)

Requirements for HANA Target Hosts and Databases

This section describes user privileges and other requirements for HANA target hosts and databases collectively referred to as the target environment.

A **target environment** is where virtualized databases run. A target environment must have a HANA instance installed and configured. The VDB will become a tenant on the target database instance.

General Target Environment Requirements

1. The operating system of the target environment must match the staging environment.
2. The SAP HANA installation on the target environment must have the same or higher version than the staging environment.
3. There must be an operating system user with the following privileges:
 1. The Delphix Engine must be able to make an SSH connection to the target environment using the operating system user.
 2. The operating system user must have read and execute privileges on the HANA binaries installed on the target environment.
 3. The operating system user must have permission to run mount and unmount as the superuser via sudo with neither a password nor a TTY. See [Sudo Requirements for the HANA 2.0 Plugin](#).
 - Prior to the discovery, a plugin directory on the target environment must exist with the following properties:
 1. The plugin directory must be writable by the operating system user mentioned above.
 2. The plugin directory must have at least 256 MB of available storage.
 - There must be a mount point directory (for example, /mnt) that will be used as the base for mount points that are created when provisioning a VDB with the following properties:
 1. The mount point directory must be writable by the operating system user mentioned above.
 2. The mount point directory should be empty.
 - TCP/IP connectivity to and from the source environment must be configured as described in [General Network and Connectivity Requirements](#).
 - The hostname and IP address must be correctly configured.
 - HANA instance on which the VDB creation is being done should have the empty ports to host the services for the tenant database. The default port number range for tenant databases is 3<instance>40–3<instance>99. This means that the maximum number of tenant databases that can be created per instance is 20. However, you can increase this by reserving the port numbers of further instances. You do this by configuring the property [multidb] reserved_instance_numbers in the global.ini file. The default value of this property is 0. If you change the value to 1, the port numbers of one further instance are available (for example, 30040–30199 if the first instance is 00). If you change it to 2, the port numbers of two further instances are available (for example, 30040–30299 if the first instance is 00). And so on.

Related Topics

- [HANA Support and Requirements](#)

Network and Connectivity Requirements for HANA Environments

Please see [General Network and Connectivity Requirements](#) and [Architecture Best Practices for Network Configuration](#) for detailed network and connectivity requirements for the Delphix Engine, including connection requirements, port allocation, and firewall and Intrusion Detection System (IDS) considerations.

Related Topics

- [Provisioning HANA VDBs](#)

Sudo Requirements for the HANA 2.0 Plugin

The following sudo privileges are requirements for the HANA 2.0 Plugin.

Privilege	Source	Target	Rationale
mkdir/rmdir	Not Required	Required	Delphix dynamically makes and removes directories under the provisioning directory during VDB operations. This privilege is optional, provided the provisioning directory permissions allow the delphix os user to make and remove directories.
mount/umount	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during VDB operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for superuser.

It is required to specify the NOPASSWD qualifier within the "sudo" configuration file. This ensures that the "sudo" command does not demand the entry of a password.

Example with user hdbadm

```
Defaults:hdbadm !requiretty
hdbadm ALL=NOPASSWD:/sbin/mount, /sbin/umount, /bin/mkdir, /bin/rmdir
```

Example of Limiting `sudo` Access for the Delphix OS User

In situations where security requirements prohibit giving the Delphix user root privileges to mount, unmount, make directory, and remove directory on the global level, it is possible to configure the `sudoers` file to provide these privileges only on specific mount points or from specific Delphix Engines, as shown in the below example.

```
Defaults:hdbadm !requiretty
hdbadm ALL=(root) NOPASSWD: \
/bin/mount *<delphix-server-name>* /mnt/provision/*, \
/bin/umount /mnt/provision/*, \
/bin/mkdir * /mnt/*
```

 Note that the following example is for illustrative purposes and the sudo file configuration options are subject to change.

Example 1

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/hana`.

Note that wildcards are allowed for the options on `mount` and `umount` because those commands expect a fixed number of arguments after the options. The option wildcard on the `mount` command also makes it possible to specify the file-system being mounted from the Delphix Engine.

However, wildcards are not acceptable on `mkdir` and `rmdir` because they can have any number of arguments after the options. For those commands, you must specify the exact options (`-p`, `-p -m 755`) used by the Delphix Engine.

Example /etc/sudoers File Configuration on the Target Environment for sudo Privileges on the VDB Mount Directory Only (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount * /hana/*, \
/bin/umount * /hana/*, \
/bin/umount /hana/*, \
/bin/mkdir -p /hana/*, \
/bin/mkdir -p -m 755 /hana/*, \
/bin/mkdir /hana/*, \
/bin/rmdir /hana/*
```

Example 2

This example restricts the **delphix_os** user's use of `sudo` privileges to the directory `/hana`, restricts the mount commands to a specific Delphix Engine hostname and IP, and does not allow user-specified options for the `umount` command.

This configuration is more secure, but there is a tradeoff with deployment simplicity. This approach would require a different sudo configuration for targets configured for different Delphix Engines.

A Second Example of Configuring the /etc/sudoers File on the Target Environment for Privileges on the VDB Mount Directory Only, and Allows Mounting Only from a Single Server (Linux OS)

```
Defaults:delphix_os !requiretty
delphix_os ALL=(root) NOPASSWD: \
/bin/mount <delphix-server-name>* /hana/*, \
/bin/mount * <delphix-server-name>* /hana/*, \
/bin/mount <delphix-server-ip>* /hana/*, \
/bin/mount * <delphix-server-ip>* /hana/*, \
/bin/mount "", \
/bin/umount /hana/*, \
/bin/umount * /hana/*, \
/bin/mkdir [*] /hana/*, \
/bin/mkdir /hana/*, \
```

```
/bin/mkdir -p /hana/*, \  
/bin/mkdir -p -m 755 /hana/*, \  
/bin/rmdir /hana/*
```

Related Topics

- [HANA Support and Requirements](#)

Managing HANA Environments and Hosts

This section covers the following topics:

- [Workflow for HANA Environments](#)
- [Adding a HANA Environment](#)
- [Editing HANA Environment Details](#)
- [Changing the Hostname or IP Address for HANA Staging and Target Environments](#)
- [Steps to modify HANA dSource Properties](#)

Workflow for HANA Environments

A staging environment must be added to the Delphix Engine before linking. Similarly, a compatible target environment must be added to the Delphix Engine in order to provide a location to provision to.

Installation of the HANA Plugin is required before the discovery of a HANA environment can occur.

Once an environment is added to the Delphix Engine, environment discovery takes place. Environment discovery is a process that enables the HANA Plugin to determine HANA installation details on a host. The same Environment Discovery process can be repeated during an Environment Refresh in order to detect new HANA installations.

Related Topics

- [Managing HANA Environments and Hosts](#)

Adding a HANA Environment

This topic describes how to add a HANA environment.

Prerequisites

- Make sure that the HANA environment in question meets the requirements described in [Requirements for HANA Staging Hosts and Databases](#) and [Requirements for HANA Target Hosts and Databases](#).

Procedure

- Login to the **Delphix Management** application.
- Click **Manage**.
- Select **Environments**.
- Next to **Environments**, click the **Actions (...) menu** and select **Add Environment**.
- In the Add Environment dialog, In the **Host OS** select **Unix/Linux** on the menu.
- In **Server Type**, select **Standalone Host**.
- Click **Next**.
- Enter **Name** for the environment.
- Enter the **Host IP address** or **fully qualified hostname**.
- Enter the **SSH port**.
The default value is 22.
- Enter an **OS Username** for the environment which should be '<sidadm>'.

 • Select **Login Type**.

 Currently, the Enterprise Password Vault option is not supported.

- Username and Password - enter the OS username and password
 - Username and Public Key - enter the OS username.
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:


1. Select **Public Key** for the **Login Type**.
 2. Click **View Public Key**.
 3. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - a. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - b. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - c. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

1. For Password, enter the password associated with the user in step 11.
 2. Select **Public Key** for the **Login Type**.

Datasets- 1656

3. Click **View Public Key**.
4. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - a. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - b. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - c. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.
5. For Password Login, click 'Validate' to test the username and password.
6. Enter a **Plugin Path**. The location specified must be a string without spaces. Ensure this directory path exists on the host in question. Environment discovery will fail if the directory does not exist on the host.
7. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
8. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**. For more information, see [Providing Your Own Oracle Java](#).
9. Click **Submit**. On success, the new environment will be added to the list in the Environments tab.
10. Select the "Databases" tab to view discovered HANA databases:



The screenshot displays the 'Databases' tab for a HANA installation. The installation details are as follows:

INSTALLATION DETAILS	DATABASES
<p>Allow Provisioning Yes</p> <p>System Administrator User hdbadm</p> <p>SAP Mount Path /hana/shared</p> <p>HANA Version 2.00.052.00.1599235305</p> <p>Hostname source</p> <p>HANA Instance Number 00</p> <p>HANA Installation Path /usr/sap/HDB/HDB00</p> <p>HANA SID HDB</p> <p>Repository Name HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00</p>	<p>No databases found on installation.</p>

Details
Databases

Installations

Filter: none ▼

Add Dataset Home

▼
HANA 2.00.052.00.1599235305 /usr/sap/HDB/HDB00

INSTALLATION DETAILS ✎

Allow Provisioning
Yes

System Administrator User
hdbadm

SAP Mount Path
/hana/shared

HANA Version
2.00.052.00.1599235305

Hostname
source

HANA Instance Number
00

HANA Installation Path
/usr/sap/HDB/HDB00

HANA SID
HDB

Repository Name
HANA 2.00.052.00.1599235305
/usr/sap/HDB/HDB00

DATABASES +

☰
HDB
✎

Allow Linking
Enabled ✎

Source Config Name (Required)
HDB

[Add dSource](#) 🗑️

Post-Requisites

To view information about an environment after you have created it:

1. Click **Manage**.
2. Select **Environments**.
3. Select the **environment name**.

Related Topics

- [Managing HANA Environments and Hosts](#)

Editing HANA Environment Details

This topic describes how to edit the details of an environment. Once an environment has been added it is possible to change any of the following values:

- Environment OS user
- Host IP Address/Name
- SSH port
- Path to plugin location

Procedure

1. Login to the **Delphix Management** application with admin credentials or as the owner of an environment.
2. Click **Manage**.
3. Select **Environments**.
4. In the **Environments** panel, click the name of an environment to view its details.
5. Next to **Environment Detail**, click the **Pencil** icon to edit.
6. Click the **Check** icon to save your edits.

Related Topics

- [Managing HANA Environments and Hosts](#)

Changing the Hostname or IP Address for HANA Staging and Target Environments

Below is the procedure to change the Host Address of the HANA Staging and Target Environments (If the Host address of the underlying staging and target server has also changed):

Procedure

1. Disable the dSources.
2. Disable the VDBs.
3. Login to the Delphix Management application.
4. Click **Manage**.
5. Select **Environments**.
6. Click on the existing environment name you want to modify and open the environment information screen.
7. In the Attributes section, click the **Pencil** icon located Next to it.
8. Update the Host address to the new one and save the results by clicking on .
9. The New Host address will now be updated.
10. Re-enable the dSources.
11. Re-enable the VDBs.

Related Links

- [Workflow for HANA Environments](#)
- [Adding a HANA Environment](#)

Steps to modify HANA dSource Properties

Select the dSource and then the dSource **Configuration** tab to modify dSource properties. For example, to change the **Source SystemDB Username** or **Password**, select the **Dataset, Configuration**, and then the **Custom** tab as shown below:

The screenshot shows the Delphix Management interface. On the left, a sidebar lists 'Datasets' with a filter set to 'none' and options to 'Expand All' or 'Collapse All'. Under 'DSOURCE', 'DSOURCE_1' (dSource - Active) is selected. The main area shows the configuration for 'DSOURCE_1' with tabs for 'Timeflow', 'Status', and 'Configuration'. Under 'Configuration', there are sub-tabs for 'Source', 'Policies', 'Masking', 'Hooks', and 'Custom'. The 'Custom' tab is active, displaying the following properties:

- NFS Mount Location**: /mnt/provision/DSOURCE_1
- SAP HANA Staging Tenant Database Name**: DSOURCE_1
- Staging Database Parameters**
 - Staging SystemDB User Name**: SYSTEM
 - Staging SystemDB Password**: *****
 - Staging Tenant Database ADMIN User Name**: SYSTEM
 - Staging Tenant Database ADMIN Password**: *****
- Customer Provided Data Backup Location**: /hana/shared/HDB/CustomBackup/BACKUP
- Customer Provided Log Backup Location**: Unset

Related Links

- [Linking HANA Data Sources](#)
- [Advanced Data Management Settings for HANA dSources](#)

Linking Data Sources and Syncing Data with HANA

This section covers the following topics:

- [HANA dSource Ingestion](#)
- [Linking HANA Data Sources](#)
- [Staging Push Implementation Details](#)
- [Advanced Data Management Settings for HANA dSources](#)
- [Working with HANA dSource Snapshots](#)

HANA dSource Ingestion

This release of the HANA 2.0 plugin uses Native Backups/Third-party backups to stage a dSource.

An External Backup is a form of HANA Native Backup that was created by users outside of Delphix. The necessary backup files must exist prior to the dSource operation and must be prepared by the user prior to ingestion by Delphix. In particular, a Delphix-supplied script must be run to persist vital backup metadata to Delphix during the dSource creation process. The backup files must be created on the source HANA instance must be made visible to the Staging host where the dSource is being created (ex: via shared storage, scp, etc).

A new data ingestion mechanism, Staging Push has been introduced that will help users to push data into the Delphix provided mount point on their own.

There are two mechanisms to create dSource on the staging server.

- Pull architecture, where the HANA plugin pulls the data into the Delphix provided mount location.
- Staging Push architecture, where the user pushes the data into the Delphix provided mount point on their own.

Related Topics

- [Linking Data Sources and Syncing Data with HANA](#)

Linking HANA Data Sources

This topic describes basic concepts behind the creation of a dSource for a HANA database using conventional pull and Staging push mechanisms.

⚠️ For Staging Push, only native/third-party backups are used to populate data in the staging database.


For Third-party Backups, users need to provide a “Backup Path” containing FULL and Incremental backups. These files will be copied to the Delphix Engine during the dSource linking process. In this scenario, the Delphix Engine will not require any credentials of the source database. Although credentials for the tenant database (for which backup has been placed) and systemdb of staging instances would be required.

Prerequisites for Customer provided Backups Ingestion Mechanism

- Users need to run the create_backup_metadata.sh. This script will create the backupTimeInfo file at the same location where backups are kept. This script will gather all the required metadata information and will store that information in the backupTimeInfo file. This information will then get used by the plugin at the time of creating dSource.
- More details regarding the script have been included in the [HANA 2.0 Plugin Tools](#).

Procedure

1. Login to the Delphix Management application.
2. Select **Manage > Environments**. Then select the Datasets tab.



3. Create a source configuration on which to build the dSource. You can do this by clicking on the + symbol. The Source configuration can have any name.
4. In the **Add dSource** wizard, select the required source configuration.
5. After selecting the source, below are the parameters required to create the dSource.
 - a. Parameters required for the conventional **pull** mechanism:
 - i. Mount Location on Staging Host
 - ii. SAP HANA Source Tenant Database Name
 - iii. Source SystemDB User Name
 - iv. Source SystemDB Password
 - v. Source Tenant Database User Name
 - vi. Source Tenant Database Password
 - vii. Customer Provided Data Backup Location: This represents the path to the location of the customer-provided backup files.

- viii. **Customer Provided Log Backup Location:** This represents the path to log the location of the customer-provided log backup files. This parameter is optional.

The screenshot shows the 'Add dSource' configuration window. On the left, a navigation pane lists steps: Preparation, Source (selected), dSource Configuration, Data Management, Policies, Hooks, and Summary. The main area is titled 'Source' and contains the following fields:

- Data Source:** A dropdown menu with a filter set to 'none'. The list includes 'test', 'staging', and 'AppData'.
- Data Type:** 'AppData'.
- NFS Mount Location:** A text input field.
- Mount Location on Staging Host:** A text input field.
- SAP HANA Staging Tenant Database Name:** A text input field with a note: 'The name of the Tenant Database to be created'.
- Use Staging Push:** An unchecked checkbox. Below it, a note states: 'Selecting this option will allow the user to manage staging database. Staging database parameters, Data Backup and Log Backup location are not required if Staging Push is enabled'.
- Staging Database Parameters:** A section with the note 'Database parameters are required' and a '+ Add' button.
- Customer Provided Data Backup Location:** A text input field with a note: 'Please provide the path to the data backup files on the staging host'.
- Customer Provided Log Backup Location:** A text input field with a note: 'Please provide the path to the log backup files on the staging host. Please use this only if logs are to be used.'

At the bottom right, there are buttons for 'Cancel', 'Back', 'Next', and 'Submit'.

- b. Parameters required for the **Staging push** mechanism
- Mount Location on Staging Host
 - SAP HANA Source Tenant Database Name
 - Enable the **Staging Push** feature by selecting the checkbox "**Use Staging Push**". With this feature, the user handles the Source Tenant database and its state on their own and you don't need to specify any other staging database parameters, data backup, and log backup locations.
- If this check box is selected, the HANA plugin will not invoke the tenant creation and recovery operation and if this checkbox is not selected then the default pull approach will be used.

This screenshot is identical to the previous one, but the 'Use Staging Push' checkbox is now checked.

- c. Click **Next**.

- Enter a name and select a group for your dSource.
Adding a dSource to a dataset group lets you set Delphix Domain user permissions for that database and its objects, such as snapshots. See the topics under [Users and Groups](#) for more information.
- Configure the SnapSync and Retention Policy. The SnapSync policy can be set to 'None' for customers who wish not to run a scheduled SnapSync job.
- Optionally, users can provide the bash script as a plugin-defined hook. Both Pre-Snapshot and Post-SnapShot hooks can be used.

9. Select **Next**.

10. At the end of the workflow, review the dSource Configuration in the Summary pane as shown below:

Add dSource

Source
dSource Configuration
Data Management
Policies
Hooks
Summary

Summary
Review the configuration profile for this dSource.

SOURCE
Data Type
AppData
Source Environment
staged
Environment User
hdbadm

DSOURCE CONFIGURATION
dSource Name
NDF
Notes
None
Dataset Group
Untitled

HOOKS
Pre Sync
None
Post Sync
None

POLICIES
Retention Policy
Default Retention

DATA MANAGEMENT
Staging Environment
staged
User
hdbadm
Restore Full Backup
True

SETTINGS
NFS Mount Location
/mnt/provision/TEST
SAP HANA Staging Tenant Database Name
NDF
Staging Database Parameters
Staging SystemDB User Name
SYSTEM
Staging SystemDB Password

Staging Tenant Database ADMIN User Name
SYSTEM
Staging Tenant Database ADMIN Password

Cancel Back Next **Submit**

Once the dSource linking job is submitted it is possible to monitor progress by selecting **Active Jobs** in the top menu bar. On successful completion, the new dSource will be visible in the list of Datasets under its assigned group.

After dSource creation using the Staging Push mechanism, an empty snapshot will be created but users can not create VDB's from the same. Also, before creating the subsequent snapshots, users will have to create the database and point the database to use the Delphix provided mount point.

Related Topics

- [Linking Data Sources and Syncing Data with HANA](#)

Staging Push Implementation Details

This topic provides implementation details of the Staging Push for the HANA plugin. Staging Push will help users to push data into the Delphix provided mount point on their own.

Staging Push gives users control over some Staging DB processes. It will give control of the staging database to the user to pull from any backup provider. Staging database files will be stored on Delphix Storage. Delphix will still be the one snapshotting the underlying data files, and gathering the metadata required to provision from the snapshot.

The below steps show how to create a dSource using the Staging Push mechanism.

dSource Creation

The first snapshot created as a part of dSource creation contains data and log directories within the Delphix created mount point. These directories are empty for the first snapshot.

The following snippet shows how the file system appears inside the Delphix created mount point.

```
hdbadm@target:/delphix/zfs> tree
.
├── data
└── log

2 directories, 0 files
hdbadm@target:/delphix/zfs> pwd
/delphix/zfs
```

 If a system user attempts to create a VDB from this snapshot, it will fail with the below error message.

Provision virtual database "N5RQVO". ✕

Error
Data volumes are not found under the data directory. Please check that the dSource snapshot contains the valid data and log volumes required for VDB creation.

Error Code
exception.plugin.user.error.detected

Suggested Action
Please take a new snapshot of the dSource with valid data and log volumes and try again.

Subsequent Snapshots

The system user needs to follow the steps and commands below before taking any subsequent snapshots.

 The system user needs to provide the values for all the parameters provided within {}.

1. Create a database in "no start" mode. Here, "no start" implies that the database is not started.

```
create database {DATABASE_NAME} SYSTEM USER {DATABASE_PASSWORD} no start.
```

The parameters used in the below example are:

- a. DATABASE_NAME - Refers to the database name to be used for staging Push. This is the same name that was provided by the system user on the UI while creating the dSource.
- b. DATABASE_PASSWORD - Refers to the password to be used for the staged database.

Example:

```
hdbsql SYSTEMDB=> create database TEST SYSTEM USER PASSWORD Delphix_123 no start;
0 rows affected (overall time 20.534486 sec; server time 20.531628 sec)
```

2. Set the HANA database to use the Delphix created mount point to store the data and log volume files. This is done by setting the `basepath_datavolumes` and `basepath_logvolumes` to `data` and `log` directory under the mount point, respectively. All the tenant data volumes and log volumes reside in the single mount point provided by Delphix.

```
//This query sets the value of basepath_datavolumes to within the delphix
storage.
alter system alter configuration('global.ini', 'DATABASE','{DATABASE_NAME}')
SET ('persistence','basepath_datavolumes') = '{DELPHIX_MOUNT_LOCATION}/data'
```

```
//This query sets the value of basepath_logvolumes to within the delphix
storage.
alter system alter configuration('global.ini', 'DATABASE','{DATABASE_NAME}')
SET ('persistence','basepath_logvolumes') = '{DELPHIX_MOUNT_LOCATION}/log'
```

Example:

```
hdbsql SYSTEMDB=> alter system alter configuration ('global.ini','DATABASE','TEST') SET ('persistence',
'basepath_datavolumes') = '/delphix/zfs/data';
0 rows affected (overall time 5720 usec; server time 3692 usec)

hdbsql SYSTEMDB=> alter system alter configuration ('global.ini','DATABASE','TEST') SET ('persistence',
'basepath_logvolumes') = '/delphix/zfs/log';
0 rows affected (overall time 5090 usec; server time 3114 usec)
```

Once the `basepath_datavolumes` and `basepath_logvolumes` parameters are set for the tenant, then their values are reflected in the `global.ini` file in the tenant database.

The location of the `global.ini` file is `<sapmnt>/<SID>/SYS/global/hdb/custom/config/DB_<tenantdbname>`.

Example:

```
# global.ini last modified 2021-12-01 12:25:11.693519 by hdbnameserver
[persistence]
basepath_datavolumes = /delphix/zfs/data
basepath_logvolumes = /delphix/zfs/log
```

3. Create a backup catalog. This is used by HANA to determine whether recovery is possible and which backups are to be used to recover the database.

The parameters used in the below example are:

- a. logDirs - Refers to the location where logs are present for the database. If the system user does not intend to use logs then an empty directory can be created and the location of the empty directory can be provided here.
- b. dataDir - Refers to the location where backups are present.
- c. -d - Refers to the location where the catalog backup will be created.

Example:

```
hdbadm@target:/usr/sap/HDB/HDB00> hdbbackupdiag --generate --logDirs '/usr/sap/HDB/HDB00/log' --dataDir '/usr/sap/HDB/HDB00/source_backup' -d '/usr/sap/HDB/HDB00/source_backup'
Generate backup catalog /usr/sap/HDB/HDB00/source_backup/log_backup_0_0_0.1631694109458 from data dirs: /usr/sap/HDB/HDB00/source_backup and log dirs: /usr/sap/HDB/HDB00/log
```

4. Perform recovery on the tenant database. This can be done in the following two ways:
 - a. Recovery using data backups only

The following parameters for the recovery command are listed below:

- i. Recover database for {DATABASE_NAME} - Represents the name of the database for which recovery is to be performed.
- ii. Until timestamp '{TIMESTAMP}' - Represents the timestamp value till which recovery is to be performed.
- iii. Clear log using catalog path ('{CATALOG_PATH}') - Refers to the path where the catalog backups are kept.
- iv. Using data path ('{DATA_PATH}') - Refers to the path where the data backups are kept.

Example:

```
hdbsql SYSTEMDB=> recover database for TEST until timestamp '2021-09-16 08:41:50.311000000' clear log using catalog path ('/usr/sap/HDB/HDB00/source_backup') using data path ('/usr/sap/HDB/HDB00/source_backup') ignore workergroups;
0 rows affected (overall time 63.235920 sec; server time 63.232622 sec)
```

- b. Recovery using both data and log backup

Along with the above parameters, you need to specify the below parameter:

- i. Using log path ('{LOCATION_OF_THE_LOG_PATH}') - Refers to the location of log backups.

Once the recovery is done, data will be available inside the Delphix mount point and the database will be in an active state.

The below snippet shows an example file system after the recovery.

```
hdbadm@target:/delphix/zfs> tree
.
├── data
│   └── mnt00001
│       ├── hdb00002.00004
│       │   ├── datavolume_0000.dat
│       │   └── __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
│       └── hdb00003.00004
│           ├── datavolume_0000.dat
│           └── __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
├── log
│   └── mnt00001
│       ├── hdb00002.00004
│       │   ├── __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
│       │   ├── logsegment_000_00000000.dat
│       │   ├── logsegment_000_00000001.dat
│       │   └── logsegment_000_directory.dat
│       └── hdb00003.00004
│           ├── __DO_NOT_TOUCH_FILES_IN_THIS_DIRECTORY__
│           ├── logsegment_000_00000000.dat
│           ├── logsegment_000_00000001.dat
│           └── logsegment_000_directory.dat
8 directories, 12 files
```

You can now proceed to take a snapshot and the VDB creation using the new snapshot.

Advanced Data Management Settings for HANA dSources

This topic describes advanced data management settings for dSources.

When linking a dSource, you can use custom data management settings to improve overall performance and match the needs of your specific server and data environment. If no specific settings are required, leverage default data management settings.

Accessing Data Management Settings

There are three ways to set or modify data management settings for dSources:

1. During the dSource linking process, in the Policies tab of the Add dSource wizard select the **Retention Policy**.
2. Or select **Manage**, then select **Policies**.
3. In the Policies screen select the **Retention** tab. To manage retention on target after the dSource snapshots have been deleted from the replication source, select the **Replica Retention** tab.
4. To add a new retention policy click the **+Retention** button, and to add a new replica retention policy on the replication target, click the **+Replica Retention** button.
5. This will open the Retention or the Replica Retention Policy wizard, depending upon the above selection. Enter a policy name along with required retention details and click **Submit**.
For more information, see [Policies for Scheduled Jobs](#).

Retention Policies - Retention/Replica Retention

Retention policies define the length of time Delphix Engine retains snapshots within its storage. To support longer retention times, you may need to allocate more storage to the Delphix Engine. The retention policy – in combination with the SnapSync policy – can have a significant impact on the performance and storage consumption of the Delphix Engine.

Replica Retention policies define how long the snapshots are retained on replicated namespaces for dSources and VDBs after they have been deleted on the replication source. Normally, the snapshots that have been deleted on the replication source engine are also deleted on the replication target engine. A new retention policy is introduced to provide an extended lifetime of such snapshots on the replication target.

Benefits of Longer Retention

Increased retention times for snapshots extends the quantity of historical data the Delphix Engine will retain.

Common use cases for longer retention include:

- SOX compliance
- Frequent application changes and development
- Caution and controlled progression of data
- Reduction of project risk
- Restore to older snapshots

Related Topics

- [Linking HANA Data Sources](#)
- [Policies for Scheduled Jobs](#)

Working with HANA dSource Snapshots

Taking a snapshot creates a new snapshot entry in the dSource's Timeflow.

Snapshot (Default)


This section lists the steps to take a snapshot and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog, select **Perform Snapshot**.
6. Click **View:** under **Timeflow** to verify the Snapshot you just created.
7. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
8. In the Delete Snapshot dialog, select **Delete**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Snapshot with Parameters

This section lists the steps to take a snapshot with parameters and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the arrow next to the Camera icon and select **Snapshot with Params...**
5. Select **Refresh Prior to Snapshot** to refresh the staging database. This will refresh the staging database prior to taking the snapshot.



Snapshot ✕

Refresh Prior to Snapshot

Refresh the state of staging database prior to taking snapshot to change its data contents. Uncheck the box if you wish to snapshot the staging database in its current state.

Cancel Perform Snapshot

6. In the Snapshot dialog, select **Perform Snapshot**.
7. Click **View:** under **Timeflow** to verify the Snapshot you just created.
8. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Provisioning and Managing VDBs from HANA dSources

This section covers the following topics:

- [Provisioning HANA VDBs: An Overview](#)
- [Provisioning HANA VDBs](#)
- [Working with HANA VDB Snapshots](#)
- [Failure During Provision](#)

Provisioning HANA VDBs: An Overview

This topic describes the basic concepts involved with provisioning of VDBs in HANA.

Provisioning of a VDB results in creation of a new tenant database in the specified HANA target environment, with storage for the new tenant backed by Delphix storage. The data inside the tenant reflects the same data in that exists in the dSource snapshot from which the VDB is being provisioned. During the provision process, Delphix will create a filesystem in Delphix specifically for the new VDB, populate the filesystem with the appropriate data files and datafile contents, then share and mount that filesystem onto the specified HANA target host. The new HANA tenant will have its storage paths re-configured to use a single subdirectory specific to the tenant, then traverse an NFS mount which points to the filesystem on Delphix for that tenant. The VDB provision operation is very fast because Delphix does not move nor copy data to create and populate the VDB filesystem.

During the VDB creation, the location of the basepath_datavolumes and basepath_logvolumes will be changed to use the storage inside the Delphix Engine. This will mean that any new services (even if a user creates a service post VDB creation by the plugin) which will now be created for the virtualized tenant will also be created on the tenant-specific filesystem on the Delphix Engine.

Setting Log Mode Overwrite

As a part of the VDB creation process, if the correct tenant username and password are provided, then the log mode for the virtualized tenant will be set to 'overwrite'. It means that "No log backups" are created. When savepoints are written, log segments are immediately freed to be overwritten by new log entries.

HANA Port Control

Starting Delphix Engine 6.0.10.0 and HANA vSDK plugin, Delphix introduces a solution that keeps the port numbers consistent throughout the VDB life cycle so that the connections made to the VDBs are not disrupted during their life cycle.

Need for User Specified Ports?

Until now, Delphix HANA was dependent on the HANA-provided ports for any VDB. Every VDB has its own ports and connections for internal and external communication. When operations such as rewind, refresh, or enable are performed on VDBs, this can cause port instability or inconsistency problems because the port numbers could change even when the VDB name remains the same. The below points further elaborates on this limitation.

- The plugin uses HANA-provided port numbers which are selected from the available pool of port numbers in the bottom-up approach that means an idle port number at the bottom will be picked first for use.
- The plugin does not have any mechanism to remember these port numbers and use them from the pool at the time of creating any tenant database again at the time of enabling, refresh, or rewind.
 - So if we had a tenant TEST having a service indexserver at 30043 port number, then after the refresh, rewind or enable, there will be a chance that
 - a. The same port number(30043) is not available
 - b. Any port number lesser than the one being used earlier(30043) is now available.

So, there are chances that the tenant TEST is provided with another port number for the same service by HANA and if there was any connection made to the tenant TEST at 30043 will now be disrupted. Hence there is a need for a solution that allows the database to keep the port number consistent throughout the VDB life cycle.

How does the HANA Port Allocation Works?

The port control feature now provides an option to specify port numbers at the time of VDB creation. This can be done in the following two ways.

User-specified HANA ports

See the procedure for [Provisioning HANA VDBs](#) where the user can specify the HANA service and port number at the time of VDB creation.

By Plugin itself

If the service and port numbers are not provided by the user on the Plugin user interface, then the plugin itself will decide the port number and will keep the ports consistent throughout the life of a VDB.

This is how the plugin chooses the ports:

- It is based on the value of the `reserved_instance_number` parameter in the `global.ini` file. This file is located at `/hana/shared/{DATABASE_SID}/global/hdb/custom/config`. If this parameter is not defined in `global.ini`, then the value of the instance number will be used.
- Ports will be decided by the HANA plugin in a top-down manner which means that the topmost port which is free will be used first and so on.

Related Topics

- [HANA Environments and Data Sources](#)
- [Provisioning and Managing VDBs from HANA dSources](#)

Provisioning HANA VDBs

This topic describes how to provision a virtual database (VDB) from a HANA dSource.

Prerequisites

- The provision requires a linked dSource built off a physical source tenant, as described in [Linking a HANA dSource](#), or a VDB.
- The HANA target environment should adhere to the [Requirements for HANA Hosts and Databases](#).
- Operating system users and database users should be configured.

VDB creation UI

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource**.
5. Select a **snapshot** from which you want to provision.

i If the first or empty dSource Snapshot that was created using the Staging Push mechanism is detected by the plugin and an error message is displayed. This check is done for every VDB creation operation.

6. Click the **Provision VDB** icon to open the Provision VDB wizard.
7. Select a target environment from the left pane.
8. The HANA 2.0 Plugin supports one installation on the same host, hence the value in the **Installation** dropdown will be already populated. Multiple drop-down options should not be available.
9. Set the **Environment User** to be the Instance Owner. Note: The picking of instance owner is only possible if you have multiple environment users set on that host.

10. Click **Next**.
11. In the **Target Configuration** pane you will need to specify the following:

 ****** indicates the mandatory fields.

- a. **NFS Mount Location ***: Delphix recommends setting the trailing directory to match the VDB name to delineate mount points and support administration and troubleshooting. In the example shown, the **NFS Mount Location** is /mnt/provision/dev where dev matches the **SAP HANA Target Tenant Database Name**.
- b. **SAP HANA Target Tenant Database Name ***: This refers to the target tenant database name which will be created upon VDB creation.
- c. **Target SystemDB User Name ***: This refers to the SystemDB username in the target database.
- d. **Target SystemDB Password ***: This refers to the SystemDB password in the target database.
- e. **Target Tenant Database ADMIN User Name ***: This refers to the username of the tenant system user.
- f. **Target Tenant Database ADMIN Password ***: This refers to the password of the tenant system user.

 The SYSTEM username and password derived from the physical source database should be entered during the provision of a VDB because those credentials will be recovered during provision.

- g. **Configure Tenants Service and Port**: This refers to the list of tenant services and port numbers in the following format.

```
<service>:<port>
For example, indexserver:30043
```

You can click "**+Add**" to add more services.

-  This is not a mandatory field. If a user doesn't provide a value to this, then the plugin will proceed with the above fields and it will allocate the service and port on its own.
- If an invalid service is provided, the plugin will proceed with input, but will eventually fail while creating the service and will display an error on the Delphix Engine UI.

12. In the configuration section
 - a. Set the vFile Name or accept the default, which is a random string. Delphix strongly recommends applying a consistent naming convention. We suggest this name should match the Target Tenant Name from the previous pane. Maintaining a consistent naming convention will support administration and troubleshooting activities. In the example shown, the vFile Name “dev” matches the actual tenant name and mount point from the previous pane.
 - b. Select a **Target Group** for the VDB.
Click the **Add Dataset Group** icon to add a new group if required.
 - c. Set **Auto VDB Restart** option to automatically restart the VDB, where the target host is rebooted. Delphix recommends enabling this setting.
13. Click **Next**.
14. VDB **Snapshot Policy**. Select the Snapshot policy if an automated VDB snapshot schedule is required. Optionally, the Snapshot policy may be disabled by selecting “None”.
15. Click **Next**.
16. **Masking**. This page offers the opportunity to link the provision job to a masking job. At this time, Delphix Masking does not support HANA. Please do not use this option. “Mask this VDB” should remain deselected.
17. **Hooks**. Run a customized shell script via a hook operation at a particular moment in the provision operation.

18. Click **Next**.
19. **Summary**. Presents an opportunity to review all the information as a part of the initial provision workflow.
20. Click **Submit**.
Once a provision job is submitted it is possible to monitor its progress by selecting “Active Jobs” in the top menu bar. On successful completion, VDB will be visible in the list of Datasets under its assigned group.

Related Topics

- [Linking Data Sources and Syncing Data with HANA](#)

Working with HANA VDB Snapshots

Taking a snapshot creates a new snapshot entry in the VDB's Timeflow.


This section lists the steps to take a snapshot and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the VDB you want to perform Snapshot.
4. Click the **Camera** icon.
5. From the Snapshot dialog, select **Yes**.
6. Click **View:** under **Timeflow** to verify the Snapshot you just created.
7. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
8. In the Delete Snapshot dialog, select **Delete**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Failure During Provision

Depending on where in the provision process failure occurs, and the type of failure, extra steps may be required:

- **Force Delete** of the provision object using the User Interface - The **Force Delete** option will remove underlying filesystems shared from the Delphix Engine, irrespective of any perceived dependencies on these filesystems. If provisioning fails very early in the process, it is common for the **Force Delete** option to be the only way to remove an object. This may happen when for example, the wrong password has been entered for the target-side system database user:



- In rare cases, removal of a failed virtual tenant may be required. For example a tenant with the name “TST”.

```
SELECT * FROM M_DATABASES WHERE DATABASE_NAME='TST';
ALTER SYSTEM STOP DATABASE TST IMMEDIATE;
DROP DATABASE TST DROP BACKUPS;
```

- If the **Force Delete** option could not unmount the virtual tenant mount points, then the unmount command should be run to remove the stale mount points.

```
umount -lf /mnt/provision/example
```

Related Topics

- [Provisioning HANA VDBs](#)

HANA Hook Operations

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at /bin/sh. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Admin application and command-line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of 0. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi
exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh
"$ARG_ENVIRONMENT_VARIABLE" "second argument in double quotes" 'third argument
in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a bash binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Admin application and command-line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of 0. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi
exit 0
```

Shell Operation Tips

Using nohup

You can use the nohup command and process backgrounding from the resource in order to "detach" a process from the Delphix Engine. However, if you use nohup and process backgrounding, you MUST redirect stdout and stderr.

Unless you explicitly tell the shell to redirect stdout and stderr in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either stdout or stderr.

Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your RunCommand operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The RunExpect operation executes an Expect script on a Unix environment. The Expect utility provides a scripting language that makes it easy to automate interactions with programs that normally can only be used interactively, such as ssh. The Delphix Engine includes a platform-independent implementation of a subset of the full Expect functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Admin application and CLI to aid in debugging.

If successful, the script must exit with an exit code of 0. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an ssh session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

Related Topics

- [Hook Scripts for Automation and Customization](#)

Unstructured Files and App Data

This section covers the following topics:

- [Getting Started with Unstructured Files](#)
- [Unstructured Files Environment Requirements](#)
- [Linking Unstructured Files](#)
- [Create an Empty VDB for Unstructured Files in the Delphix Engine](#)
- [Provisioning Unstructured Files as vFiles](#)
- [Managing vFiles](#)
- [vFiles Best Practices and Common Pitfalls](#)
- [Delphix Engine Plugin Management](#)
- [Unstructured Files Hook Operation Notes](#)

Getting Started with Unstructured Files

The term “unstructured files” refers to data stored in a filesystem that is NOT usually accessed by a DBMS or similar software. Unstructured files can consist of anything from a simple directory to the root of a complex application like Oracle Enterprise Business Suite. Like with other data types, you can configure a dSource to sync periodically with a set of unstructured files external to the Delphix Engine. The dSource is a copy of these physical files stored on the Delphix Engine. On Unix platforms, dSources are created and periodically synced by an implementation of the rsync utility. On Windows, files are synced using the robocopy utility, which is distributed with Windows. dSources enable you to provision “vFiles,” which are virtual copies of data that are fully functional read-write copies of the source of the original file. You can mount vFiles across one target environment or many.

Unstructured Files Environment Requirements

This section covers the following topics:

- [Unstructured Files on Unix Environments](#)
- [Unstructured Files on Windows Environments](#)

Unstructured Files on Unix Environments

This section covers the following topics:

- [Requirements for Unix Environments](#)
- [Network and Connectivity Requirements for Unix Environments](#)
- [Sudo Privilege Requirements for Unstructured Files on Unix](#)
- [Sudo File Configuration Examples for Unstructured Files on Unix](#)
- [Adding a Unix Environment](#)

Requirements for Unix Environments

This topic outlines the supported operating systems (OSs) for use on UNIX source and target environments.

Supported Operating Systems

Operating System	Version	Processor Family
Solaris	10, 11	SPARC x86_64
Red Hat Enterprise Linux	5.5 - 5.11 6.1 - 6.10 7.0 - 7.8 8.0	x86_64
Oracle Enterprise Linux	5.5 - 5.11 6.1 - 6.10 7.0 - 7.8 8.0	x86_64
SUSE Linux Enterprise Server	11, 11SP1	x86_64
AIX	6.1, 7.1, 7.2	Power
HP-UX	11i v3 (11.31)	IA64
Ubuntu	18, 20	X86_64

Delphix supports all 64-bit OS environments for source and target.

Required HP-UX patch for Target Servers

PHNE_37851 - resolves a known bug in HP-UX NFS client prior to HP-UX 11.31.

Additional Source Environment Requirements

- There must be an operating system user with these privileges. For example, here in this section `delphix_os` is the primary user for the environment:
- Ability to login to the source environment via SSH
- There must be a directory on the source environment where you can install the Delphix platform Toolkit – for example, `/var/opt/delphix/toolkit`.
- The **delphix_os** user must own the directory
- The directory must have permissions `-rwxrwx---` (0770), but you can also use more permissive settings

- The **delphix_os** user must have read and execute permissions on each directory in the path leading to the toolkit directory. For example, when the toolkit is stored in `/var/opt/delphix/toolkit`, the permissions on `/var`, `/var/opt`, and `/var/opt/delphix` should allow read and execute for "others," such as `-rwxr-xr-x`.
- The directory should have 1.5GB of available storage: 400MB for the toolkit and 400MB for the set of logs generated by each client that runs out of the toolkit
- On a Solaris host, `gtar` must be installed. Delphix uses `gtar` to handle long file names when extracting the toolkit files into the toolkit directory on a Solaris host. The `gtar` binary should be installed in one of the following directories:
 - `/bin:/usr`
 - `/bin:/sbin:/usr`
 - `/sbin:/usr/contrib`
 - `/bin:/usr/sf`
 - `/bin:/opt/sfw`
 - `/bin:/opt/csw/bin`
- The Delphix Engine must be able to initiate an SSH connection to the source environment

Additional Target Environment Requirements

- There must be an operating system user with these privileges. Here we use `delphix_os` as an example for the primary user for the environment:
- Ability to login to the target environment via SSH
- The following permissions are usually granted via sudo authorization of commands.
- See [Sudo Privilege Requirements](#) for further explanation of the commands and [Sudo File Configuration Examples for Unstructured Files on Unix](#) for examples of the `/etc/sudoers` file on different operating systems.
 - The primary user for the target environment must have the ability to run `mount`, `umount`, `mkdir`, and `rmdir` as a super-user on any directory mounted to by the Delphix Engine.
 - If the target host is an AIX system, permission to run the `nfso` command as a super-user
- There must be a directory on the [target environment](#) where you can install the Delphix Engine Toolkit - for example, `/var/opt/delphix/toolkit`.
 - The **delphix_os** user must own the directory
 - The directory must have permissions `-rwxrwx---` (0770), but you can also use more permissive settings
 - The **delphix_os** user must have read and execute permissions on each directory in the path leading to the toolkit directory. For example, when the toolkit is stored in `/var/opt/delphix/toolkit`, the permissions on `/var`, `/var/opt`, and `/var/opt/delphix` should allow read and execute for "others," such as `-rwxr-xr-x`.
 - The directory should have a total of at least 800MB of storage, plus 1MB of storage per vFile that will be provisioned to the target
- On a Solaris host, `gtar` must be installed. Delphix uses `gtar` to handle long file names when extracting the toolkit files into the toolkit directory on a Solaris host. The `gtar` binary should be installed in one of the following directories:
 - `/bin:/usr`

- `/bin:/sbin:/usr`
- `/sbin:/usr/contrib`
- `/bin:/usr/sf`
- `/bin:/opt/sfw`
- `/bin:/opt/csw/bin`
- There must be a directory under which the mount points are created when provisioning a vFile to the target environment. The group associated with the directory must be the primary group of the delphix-os user. Group permissions for the directory should allow read, write, and execute by members of the group.
- The Delphix Engine must be able to initiate an SSH connection to the target environment.
- NFS client services must be running on the target environment.
 - Required packages on target hosts:
 - i. Portmapper (rpcbind)
 - ii. status daemon (rpc.statd)
 - iii. NFS lock manager (rpc.lockd/lockmgr)
 - The Delphix Engine enables the use of NFSv3 by default for mounting target host filesystems, thus, the prerequisite packages to support NFSv3 client communication are required for normal operation. In addition, the required services to support NFS client communications (including file locking) must be running. These services are shown in the left column.
 - To enable NFSv4, which does not need to interact with those discrete services, See [NFSv4 Configuration](#).

Related Topics

- [Unstructured Files on Unix Environments](#)
- [Sudo File Configuration Examples for Unstructured Files on Unix](#)
- [Sudo Privilege Requirements for Unstructured Files on Unix](#)

Network and Connectivity Requirements for Unix Environments

Port Allocations Specific to Unstructured Files

The Delphix Engine makes use of the following network ports for unstructured files dSources and vFiles:

Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	873	Rsync connections used for communication to <code>rsyncd</code> during SnapSync
TCP/UDP	111	Remote Procedure Call (RPC) port mapper used for NFSv3 mounts Note: RPC calls in NFSv3 use additional fixed ports for supporting services (lockd, mountd and statd) seen below.
TCP	1110	NFS Server daemon status and NFS server daemon keep-alive (client info)
TCP	2049	NFS Server daemon from vFiles to the Delphix Engine
TCP	54043	NFSv3 mount daemon
TCP	54044	NFSv3 stat daemon (lock state notification service)
TCP	54045	NFSv3 lock daemon/manager
UDP	33434 - 33464	Traceroute from source and target database servers to the Delphix Engine (optional)

Outbound from a Source Environment Port Allocation

Protocol	Port Numbers	Use
TCP	873	Rsync connections used during SnapSync
TCP	xxxx	DSP connections used for monitoring and script management during SnapSync. Typically DSP runs on port 8415.

Inbound to a Source Environment Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to the source environment

Outbound from a Target Environment Port Allocation

Protocol	Port Numbers	Use
TCP	873	Rsync connections used during V2P
TCP	xxxx	DSP connections used for monitoring and script management. Typically DSP runs on port 8415.

Inbound to a Target Environment Port Allocation

Protocol	Port Numbers	Use
TCP	22	SSH connections to the target environment

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

Related Topics

- [Unstructured Files on Unix Environments](#)
- [Requirements for Unix Environments](#)

Sudo Privilege Requirements for Unstructured Files on Unix

This topic describes `sudo` file privilege configurations necessary for interacting with the Delphix Engine when virtualizing unstructured files on Unix Environments.

Sudo Privilege Rationale

Privilege	Sources	Targets	Rationale
<code>mkdir/rmdir</code>	Not Required	Optional	Delphix dynamically makes and removes directories under the provisioning directory during vFiles operations. This privilege is optional, provided the provisioning directory permissions allow the delphix_os user to make and remove directories.
<code>mount/umount</code>	Not Required	Required	Delphix dynamically mounts and unmounts directories under the provisioning directory during vFiles operations. This privilege is required because <code>mount</code> and <code>umount</code> are typically reserved for a super-user.
<code>nfso</code> (AIX only)	Not Required	Required	Delphix monitors NFS read and write sizes on an AIX target host. It uses the <code>nfso</code> command to query the sizes in order to optimize NFS performance for vFiles running on the target host. Only a super-user can issue the <code>nfso</code> command.

Related Topics

- [Unstructured Files on Unix Environments](#)
- [Requirements for Unix Environments](#)

Sudo File Configuration Examples for Unstructured Files on Unix

This topic describes `sudo` file privilege configurations necessary for interacting with the Delphix Engine when virtualizing unstructured files on Unix Environments.

Considerations for sudo access and account locking

The Delphix Engine tests its ability to run the `mount` command using `sudo` on the target environment by issuing the `sudo mount` command with no arguments. Many of the examples shown in this topic do not allow that. This causes a warning during environment discovery and monitoring but otherwise does not cause a problem. If your vFiles operations succeed, it is safe to ignore this warning.

However, some users configure the security on the target environments to monitor `sudo` failures and lockout the offending account after some threshold. In those situations, the failure of the `sudo` commands might cause the `delphix_os` account to become locked. One workaround for this situation is to increase the threshold for locking out the user account. Another option is to modify `/etc/sudoers` to permit the `delphix_os` user to run the `mount` command without parameters.

Configuring `sudo` Access on Solaris for Unstructured Files

On a Solaris SPARC target, `sudo` access to `mount`, `umount`, `mkdir`, and `rmdir` is required. In this customer example, super-user privileges are restricted to the virtual dataset mount directory `/delphix` and are further restricted to commands which mount data from a single Delphix Engine with IP address `100.245.235.12`.

Delphix requires `umount -f` for emergency force unmounts on Solaris.

Additionally, `sudo` access to `ps` may be added to avoid warnings during discovery but is not required.

Example: Solaris `/etc/sudoers` entries for a Delphix Target for Unstructured Files

```
User_Alias DELPHIX_USER=delphix_os

Cmd_Alias DELPHIX_CMDS= \
/usr/sbin/mount 100.245.235.12\:* /delphix/*, \
/usr/sbin/mount * 100.245.235.12\:* /delphix/*, \
/usr/sbin/umount /delphix/*, \
/usr/sbin/umount * /delphix/*, \
/usr/sbin/umount -f /delphix/*, \
/usr/bin/mkdir /delphix/*, \
/usr/bin/mkdir -p /delphix/*, \
/usr/bin/rmdir /delphix/*
/usr/bin/ps
DELPHIX_USER ALL=(ALL) NOPASSWD: DELPHIX_CMDS
```

Configuring `sudo` Access on Linux for Unstructured Files

On a Linux target, sudo access to `mount`, `umount`, `mkdir`, and `rmdir` is required. In this customer example, super-user privilege is restricted to the virtual database mount directory `/delphix`. Aliases are used to restrict the Delphix Engines which are allowed to run these commands.

Delphix requires `umount -lf` for emergency force unmounts on Linux.

Example: Linux `/etc/sudoers` file for a Delphix Target for Unstructured Files

```
Defaults:delphix_os !requiretty

Cmd_Alias DELPHIX_ADMIN_CMDS= \
/bin/mount /delphix/*, \
/bin/mount * /delphix/*, \
/bin/umount /delphix/*, \
/bin/umount * /delphix/*, \
/bin/umount -lf /delphix/*, \
/bin/mkdir -p -m 755 /delphix/*, \
/bin/mkdir -p /delphix/*, \
/bin/mkdir /delphix/*, \
/bin/rmdir /delphix/*
/bin/ps
Host_Alias DELPHIX_HOSTS=delphix001, delphix002
delphix_os DELPHIX_HOSTS=NOPASSWD:DELPHIX_ADMIN_CMDS
```

Configuring `sudo` Access on AIX for Unstructured Files

In addition to sudo access to the `mount`, `umount`, `mkdir`, and `rmdir` commands on AIX target hosts, Delphix also requires sudo access to `nfso`. This is required on target hosts for Delphix to monitor the NFS read/write sizes configured on the AIX system. The super-user access level is needed to run the `nfso` command. This example does not restrict the Delphix Engine which is allowed to run these commands.

Delphix requires `umount -f` for emergency force unmounts on AIX.

Example: AIX `/etc/sudoers` File for a Delphix Target for Unstructured Files

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD: \
/bin/mount, \
/bin/umount, \
/bin/mkdir, \
/bin/rmdir, \
/usr/sbin/nfso, \
/usr/bin/ps
```

Configuring `sudo` Access on HP-UX for Unstructured Files

On the HP-UX target, as with other operating systems, `sudo` access to `mount`, `umount`, `mkdir`, and `rmdir` is required. This example does not restrict the Delphix Engine which are allowed to run these commands.

Example: HP-UX `/etc/sudoers` file for a Delphix Target for Unstructured Files

```
Defaults:delphix_os !requiretty
delphix_os ALL=NOPASSWD:/sbin/mount, /sbin/umount, /bin/mkdir, /bin/rmdir, /bin/ps
```

Related Topics

- [Unstructured Files on Unix Environments](#)
- [Requirements for Unix Environments](#)

Adding a Unix Environment

This topic describes how to add a new Unix environment.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Click the **Plus** icon next to **Environments**.
5. In the **Add Environment** dialog, select **Unix/Linux**.
6. Select **Standalone Host**.
7. Enter the **Host IP address**.
8. Enter an optional **Name** for the environment.
9. Enter the **SSH** port. The default value is **22**.
10. Enter a **Username** for the environment.
11. Select **Login Type**.
 - Password - enter the OS password associated with the user in Step 10, or
 - Public Key, or
 - Password Vault - select from an existing Enterprise Password Vault

Using Public Key Authentication

If you want to use public-key authentication for logging into your Unix-based environment, there are two options: use the engine's key pair or provide a key pair for this environment.

To use the engine's key pair:

- a. Select **Public Key** for the **Login Type**.
- b. Click **View Public Key**.
- c. Copy the public key that is displayed, and append it to the end of your `~/.ssh/authorized_keys` file. If this file does not exist, you will need to create it.
 - i. Run `chmod 600 ~/.ssh/authorized_keys` to allow only the file's owner to read and write to it (make sure the file is owned by the user).
 - ii. Run `chmod 755 ~` to restrict access to the user's home directory so no other user may write to it.
 - iii. Run `chmod 700 ~/.ssh` so that others cannot write to it. The `~/.ssh` directory cannot be writable by group or other users. Otherwise, authentication will fail.

As an alternative, you can provide a key pair specific for this environment via the API or CLI. See [Option 2 in this CLI Cookbook article](#) for instructions.

12. For **Password Login**, click **Verify Credentials** to test the username and password.
13. Enter a **Toolkit Path**.
The toolkit directory stores scripts used for Delphix Engine operations. It should have a persistent working directory rather than a temporary one.
14. Click **Submit**.

Post-Requisites

After you create the environment, you can view information about it by doing the following:

1. Click **Manage**.
2. Select **Environments**.
3. Select the **environment name**.

Related Topics

- [Unstructured Files on Unix Environments](#)

- [Requirements for Unix Environments](#)

Unstructured Files on Windows Environments

This section covers the following topics:

- [Requirements for Windows Environments](#)
- [Network and Connectivity Requirements for Windows Environment](#)
- [Adding a Windows Environment](#)
- [Options for Linking Unstructured Files on Windows Environments](#)

Requirements for Windows Environments

Supported Operating Systems

- Windows Server 2012, 2012 R2
- Windows Server 2016
- Windows Server 2019
- Windows Server 2022

Requires 64-Bit Windows

Delphix must install the Delphix Connector on all Windows hosts that Delphix will directly communicate with. This means all target hosts, and source or staging hosts. The Delphix Connector only supports 64-bit versions of Windows.

See [Options for Linking Unstructured Files on Windows Environments](#) for more information about source vs. staging hosts.

Additional Source or Staging Environment Requirements

- The Delphix Connector must be installed on the source or staging environment. You must have used the Delphix Connector to register this environment with the Delphix Engine.
- The `robocopy` utility must be installed on the source or staging Windows environment.
`robocopy` is installed by default on Windows Server 2008, Windows Vista, Windows 7, and Windows 8. For other versions of Windows, it is available by downloading a resource kit from Microsoft.
- If using a staging environment, the source's files must be made available and readable to the environment user from the staging environment via a UNC path. For example, use Windows Sharing.

Additional Target Environment Requirements

- The Delphix Connector must be installed on the target environment. You must have used the Delphix Connector to register this environment with the Delphix Engine.

Procedure for Adding and Installing the Delphix Connector for Windows

All Windows environments that will communicate with Delphix must have the Delphix Connector installed. The instructions in this topic cover downloading Delphix Connector, running the Delphix Connector installer on the Windows machine, and then registering the environment with the Delphix Engine.

Procedure

Downloading the Delphix Connector

-  Delphix Connector software supplied by Delphix Engine versions before 4.2.4.0 required that the Windows machine had SQL Server installed. If you are using a Windows machine that does not have SQL Server installed, you must download the Delphix Connector from a Delphix Engine of version 4.2.4.0 or higher.

The Delphix Connector can be downloaded through the Delphix Engine Interface, or by directly accessing its URL.

Using the Delphix Engine Interface

 A Flash player must be available on the Windows host in order to download Delphix Connector using the Delphix GUI.

1. From the Windows machine that you want to use, start a browser session and connect to the **Delphix Management** application using the `delphix_admin` login.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Plus** icon.
5. In the **Add Environment** dialog, select **Windows** in the operating system menu.
6. Select **Target**.
7. Select **Standalone**.
8. Click the download link for the **Delphix Connector Installer**.
The Delphix Connector will download to your local machine.

Direct Download

1. You can download the Delphix Connector directly by navigating to this URL: `http://<name of your Delphix Engine>/connector/DelphixConnectorInstaller.exe`

Installing Delphix Connector

On the Windows machine that you want to want to use, run the Delphix Connector installer. Click **Next** to advance through each of the installation wizard screens.

 The installer will only run on 64-bit Windows systems. 32-bit systems are not supported.

1. For **Connector Configuration**, make sure there is no firewall in your environment blocking traffic to the port on the Windows environment that the Delphix Connector service will listen to.
2. For **Select Installation Folder**, either accept the default folder or click **Browse** to select another.
3. Click **Next** on the installer final **Confirm Installation** dialog to complete the installation process and then **Close** to exit the Delphix Connector Install Program.
4. Note: At this point, you can close the Delphix GUI dialog by clicking **Cancel**.

Registering Environment With Delphix Engine

1. Return to the Delphix Management application.
2. Enter the **Environment Name**, **Host Address**, **Delphix Connector Port**, **OS Username**, and **OS Password** for the target environment.
3. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
4. In the Java Development Kit tab, enter the absolute path to your Oracle JDK and click **Next**.
5. Click **Submit**.

As the new environment is added, you will see two jobs running in the **Delphix Admin Job History**, one to **Create and Discover** an environment, and another to **create** an environment. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel.

Post-Requisites

- On the Windows machine, in the **Windows Start Menu**, go to **Services > Extended Services**, and make sure that the **Delphix Connector** service has a **Status** of **Started**, and that the **Startup Type** is **Automatic**.

Related Topics

- [Unstructured Files on Windows Environments](#)
- [Options for Linking Unstructured Files on Windows Environments](#)

Network and Connectivity Requirements for Windows Environment

Port Allocations Specific to Unstructured Files

The Delphix Engine makes use of the following network ports for unstructured files dSources and VDBs:

Outbound from the Delphix Engine

Protocol	Port Numbers	Use
TCP	xxxx	Delphix Connector connections to source and target environments. Typically, the Delphix Connector runs on port 9100.

Inbound to the Delphix Engine

Protocol	Port Number	Use
TCP	3260	iSCSI target daemon for connections from iSCSI initiators on the target environments to the Delphix Engine

Outbound from a Source, Staging, or Target Environment

Protocol	Port Numbers	Use
TCP	80	The Delphix Connector registers environments over HTTP
TCP	xxxx	DSP connections used for monitoring and script management. Typically, DSP runs on port 8415.

Inbound to a Source, Staging, or Target Environment

Protocol	Port Numbers	Use
TCP	xxxx	Delphix Connector connections to source environments. Typically, the Delphix Connector runs on port 9100.
TCP	445	If a staging environment is used, the staging server must be able to access SMB (Windows File Sharing) files on the source server.

General Outbound from the Delphix Engine Port Allocation

Protocol	Port Numbers	Use
TCP	25	Connection to a local SMTP server for sending email
TCP/UDP	53	Connections to local DNS servers
UDP	123	Connection to an NTP server
UDP	162	Sending SNMP TRAP messages to an SNMP Manager
TCP	443	HTTPS connections from the Delphix Engine to the Delphix Support upload server
TCP/UDP	636	Secure connections to an LDAP server
TCP	8415	Connections to a Delphix replication target. See Configuring Replication .
TCP	50001	Connections to source and target environments for network performance tests.

General Inbound to the Delphix Engine Port Allocation

Protocol	Port Number	Use
TCP	22	SSH connections to the Delphix Engine
TCP	80	HTTP connections to the Delphix GUI
UDP	161	Messages from an SNMP Manager to the Delphix Engine
TCP	443	HTTPS connections to the Delphix Management Application
TCP	8415	Delphix Session Protocol connections from all DSP-based network services including Replication, SnapSync for Oracle, V2P, and the Delphix Connector.
TCP	50001	Connections from source and target environments for network performance tests via the Delphix CLI.

Firewalls and Intrusion Detection Systems (IDS)

Production databases on source environments (for dSources) are often separated from the non-production environment by firewalls. Firewalls can add milliseconds to the latency between servers. Accordingly, for best performance, there should be no firewalls between the Delphix Engine and the virtual database (VDB) target environments. If the Delphix Engine is separated from a source environment by a firewall, the firewall must be configured to permit network connections between the Delphix Engine and the source environments for the application protocols (ports) listed above.

Intrusion detection systems (IDSs) should also be made permissive to the Delphix Engine deployment. IDSs should be made aware of the anticipated high volumes of data transfer between dSources and the Delphix Engine.

Related Topics

- [Unstructured Files on Windows Environments](#)
- [Requirements for Windows Environments](#)

Adding a Windows Environment

This topic describes how to add a Windows environment to the Delphix Engine for use with unstructured files.

All Windows source and target environments containing unstructured files must have the Delphix Connector installed to enable communication between the environment and the Delphix Engine. The instructions in this topic cover initiating the Add Target process in the Delphix Management application, running the Delphix Connector installer on the environment, and verifying that the environment has been added to the Delphix Engine.

Prerequisites

- Make sure your source and target environment meet the requirements described in [Requirements for Windows Environments](#).

Procedure

1. From the machine that you want to use, login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Next to **Environments**, click the **Actions** menu and select **Add Environment**.
5. In the **Add Environment** wizard Host and Server tab, select:
 - a. Host OS: **Windows**
 - b. Host Type: **Target**.
 - c. Server Type: **Standalone**.
6. Click **Next**.
7. In the Environment Settings tab click the download link for the **Delphix Connector Installer**. The Delphix Connector will download to your local machine.
8. On the Windows machine that you want to use as a target, run the Delphix Connector installer. Click **Next** to advance through each of the installation wizard screens. The installer will only run on 64-bit Windows systems. 32-bit systems are not supported.
 - a. For **Connector Configuration**, make sure there is no firewall in your environment blocking traffic to the port on the target environment that the Delphix Connector service will listen to.
 - b. For **Select Installation Folder**, either accept the default folder or click **Browse** to select another.
 - c. Click **Next** on the installer final 'Confirm Installation' dialog to complete the installation process and then **Close** to exit the Delphix Connector Install Program.
9. Return to the Delphix Management application.
10. Enter the **Environment Name, Host Address, Delphix Connector Port, OS Username, and OS Password** for the target environment.
11. To provide your own Oracle Java select the **Provide my own JDK** checkbox and click **Next**.
12. In the Java Development Kit tab enter the absolute path to your Oracle JDK and click **Next**.
13. Click **Submit**.

As the new environment is added, you will see two jobs running in the **Delphix Admin Job History**, one to **Create and Discover** an environment, and another to **Create** an environment. When the jobs are complete, you will see the new environment added to the list in the **Environments** panel.

Post-Requisites

1. On the Windows environment, in the **Windows Start Menu**, select **Services**.
2. Select **Extended Services**.
3. Make sure that the **Delphix Connector** service has a Status of **Started**.
4. Make sure that the **Startup Type** is **Automatic**.

Related Topics

- [Unstructured Files on Windows Environments](#)

- [Requirements for Windows Environments](#)

Options for Linking Unstructured Files on Windows Environments

There are two techniques for linking a new dSource from files on a Windows source.

Direct Communication With The Source Environment

The simplest technique is to have the Delphix Engine communicate directly with the source environment. This requires installing the Delphix Connector on the source machine. When linking, specify a local path on the source machine, such as C:\Files\MyData.

Using a Staging Environment

In some cases, it is not possible or desirable to install the Delphix Connector on the source environment. In those cases, you can install the Delphix Connector on a "staging environment." This is another Windows machine that will act as an intermediary between the Delphix Engine and the source environment. Files on the source must be accessible by the environment user from the staging environment via a UNC path. Specifically, the environment user is only required to have READ access to the path, its directories, and files so that the robocopy utility called from the staging host can function properly. For example, use Windows Sharing on the source machine. When linking, specify the UNC path to the files on the source – for example, \\MySource\MyData\

For more information on installing the Delphix connector, refer to [Installing the Delphix Connector Service on the Target Database Servers](#).

Related Topics

- [Unstructured Files on Windows Environments](#)
- [Requirements for Windows Environments](#)

Linking Unstructured Files

Prerequisites

- The source environment must meet the requirements outlined in [Unstructured Files Environment Requirements](#).
- The Delphix Engine must have access to an environment user. This user should have read permissions on all files to be cloned.

i Unstructured Files on Cluster Environments

Unstructured files cannot be linked from, or provisioned to, any form of a cluster environment, such as an Oracle RAC environment. To link or provision unstructured files from a host that is part of a cluster, add the host as a standalone environment. Then link from, or provision to, this standalone host.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **environment** containing the unstructured files you want to link.
5. Click the **Environment Details** tab.
6. If the environment user described in the Prerequisites section is not listed under **Environment Users**, add the user.
7. Click the **Databases** tab.
8. Scroll to the bottom of the page to view the **Unstructured Files** section.
9. Click the **Plus** icon on the right. This action displays a dialog box prompting for the **Database Name** and **Path**. Enter a name to help identify the files. The path is the absolute path to the directory on the environment server.
10. Click **Add** to save the configuration. After saving this configuration, add the dSource.
11. Click **Manage > Datasets**.
12. On the left-hand side, click the **plus** sign.
13. Select **Add dSource**.
Alternatively, on the **Environment Management** screen, you can click **Add dSource** next to a dataset name to start the dSource creation process.
14. In the **Add dSource** wizard, select the source of the files.
15. Select the **Environment User** outlined in the Prerequisites section.
16. Click **Advanced**.
17. Enter **Paths to Exclude**.
These paths are relative to the root path of the dataset home path and will not be linked by the Delphix Engine. This feature is most commonly used to exclude directories containing log files. Wildcard (*) pattern matching is supported to exclude all the contents of a directory, without excluding the directory itself. For example, specifying `/dir/*` will exclude all contents of `/dir` but still link `/dir` as an empty directory. For PowerShell to escape a \$ sign in a directory path please use the following `/`` before the dollar sign when adding it to the exclude paths, for example: `/` $RECYCLE.BIN`.

i Retroactive Edits to Exclude Paths on Windows

After creating a dSource, you can edit the set of **Paths to Exclude** from syncing at any time on the dSource's **Configuration** tab. For Unix environments, retroactively adding a path to exclude will result in the next SnapSync deleting the newly-excluded files. However, for Windows environments, retroactively adding a path to exclude will result in the next SnapSync ignoring newly-excluded files. Stale versions of these files will still exist in all future snapshots.

18. If you are linking files from a Unix environment, enter **Paths of Symlinks to Follow**. These paths are relative to the root path of the dataset home path and will be followed to gather additional files to copy.

i Paths of Symlinks to Follow - Caveats

- This feature can only be used to follow symlinks to directories. Symlinks to files will be ignored.
- This feature is not available for files on Windows environments.

19. Click **Next**.
20. Enter a **dSource Name**.
21. Select a **Database Group** for the dSource.
22. Click **Next**.
Adding a dSource to a database group enables you to set Delphix Domain user permissions for that dSource's objects, such as snapshots.
23. Select a **SnapSync** policy.
24. Click **Advanced** to edit retention policies.
25. Click **Next**.
26. Enter any operations that should be run at **Hooks** during the sync process (or any future sync processes).
27. Click **Next**.
28. Review the **dSource Configuration** and **Data Management** information.
29. Click **Submit**.

The Delphix Engine will initiate two jobs to create the dSource, **DB_Link**, and **DB_Sync**. You can monitor these jobs by clicking **Active Jobs** in the top menu bar, or by selecting **System > Event Viewer**. When the jobs have completed successfully, the file's icon will change to a **dSource** icon on the **Environments > Databases** screen, and the dSource will be added to the list of **Datasets** under its assigned group.

i dSource Information

After you have created a dSource, you can view information about it and make modifications to its policies and permissions by selecting it in the **Datasets** panel.

Related Topics

- [Unstructured Files and App Data](#)

Create an Empty VDB for Unstructured Files in the Delphix Engine

This topic describes the procedure for creating an empty VDB, used for unstructured files. The term "unstructured files" in Delphix refers to a dataset that acts as a directory of different files. An empty VDB for unstructured files is not a database and does not receive special treatment or processing by Delphix, it exists as a place for the files to be generated, tracked, and copied. Creating a VDB can be done with provisioning from an existing dataset (a dSource or another VDB) or as an empty VDB created and filled with data.

Creating an empty VDB places an initially-empty mount on target environments. It functions similar to a VDB created via provisioning except it cannot refresh. Refreshing a dataset means overwriting the dataset content with new data pulled from the dataset parent. If a new VDB is created from scratch, the newly-created dataset will not have a parent and cannot be refreshed. All other functions are identical, meaning that the new VDB for unstructured files can be provisioned from, rewind, take snapshots, and so on.

Prerequisites

The target environment must meet the requirements outlined in [Unstructured Files Environment Requirements](#).

Unstructured Files on Cluster Environments

You cannot create an empty VDB on any form of the cluster environment, such as an Oracle RAC environment. To create an empty VDB on a host that is part of a cluster, add the host as a standalone environment, then create the empty VDB on the standalone host.

Procedure

To create an empty VDB without provisioning:

1. Log in to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **plus** icon.
5. Select **Create Empty VDB**.
6. Click **Next** to go to the Target environment tab and do the following:
 - a. Under Environment, select a target environment where your empty VDB will be placed.
 - b. In the **Mount Path** field, enter the absolute path where the empty VDB/dataset will be mounted.
 - c. Click **Next**.
7. On the **Configuration** tab, do the following:
 - a. In the **Empty VDB name** field, enter a name for the VDB.
 - b. Select a target group for the empty VDB.
 - c. Select the **Enable auto-restart of the empty VDB** checkbox to allow VDB to be automatically restarted when the target host is rebooted.
 - d. Click **Next**.
8. On the **Policies** tab, select a Snapshot Policy for the empty VDB.
9. (Optional) On the **Hooks** tab, select a hook point and then click + to add a script to run at that point. Click **Next**.
10. The **Summary** tab will enable you to review your configurations. Click **Submit**

After the operation completes, the empty VDB will appear in the Datasets panel.

Related Topics

- [Unstructured Files and App Data](#)

Provisioning Unstructured Files as vFiles

Overview

This topic describes the process of provisioning to a set of unstructured files as vFiles.

Prerequisites

- You will need an unstructured files dSource, as described in [Linking Unstructured Files](#), or an existing vFiles from which you want to provision another.
- The target environment must meet the requirements outlined in [Unstructured Files Environment Requirements](#).

Unstructured Files on Cluster Environments

Unstructured files cannot be linked from, or provisioned to, any form of a cluster environment, such as an Oracle RAC environment. To link or provision unstructured files from a host that is part of a cluster, add the host as a standalone environment. Then, link from or provision to this standalone host.

Post Provision/Migration Ownership Rules

When a new vFile VDB is provisioned, the ownership is changed to match the Environment User anytime a VDB is enabled. This could cause conflict in the ownership of existing files in a case where the VDB has just been migrated to a new host – the VDB files on the new host will now be owned by the new Environment User.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select a **dSource** or **vFiles**.
5. Click the **TimeFlow** tab.
6. Select a **snapshot**.
7. Click **Provision**.
The **Provision vFiles** panel will open, and the field **Mount Path** will auto-populate with the path to the files on the source environment.
8. Select a target environment.
If you need to add a new target environment for the vFiles, click the **Plus** icon next to **Filter Target** to add an environment.

 You can only target a Unix environment when provisioning from a Unix dSources or vFiles. You can only target a Windows environment when provisioning from a Windows dSources and vFiles.

9. If necessary, modify the **Mount Path**.
 - On Windows, this mount path must not be a UNC path. It must be a local drive letter and folder path. The UNC path will operate after the provisioning completes.
 - On Linux and Unix hosts, this mount path must be the full path and not include symlinks.
10. Click **Advanced**.

11. Enter **Additional Mount Points**.

When it is mounted to the target environment, the vFiles will be mounted to any additional mount points you provide.

 The **Shared Path** is a relative path dictating which portion of the vFiles should be available on the additional environments. To share the entirety of the vFiles, specify a **Shared Path** of /.

12. Select an Environment User to own the mounted files.

If the files are being mounted to multiple environments, ensure this user is available across all environments.

13. Click **Next**.

14. Enter a vFiles **Name**.

15. Select a **Target Group** for the vFiles.

If necessary, click the **Plus** icon to add a new group.

16. Select a **Snapshot Policy** for the vFiles. if necessary create a new policy.

17. Click **Next**.

18. Enable Auto VDB Restart to allow the VDB to be automatically restarted when the target host reboot is detected by Delphix.

19. Enter any operations that should be run as Hooks during the lifetime of the vFiles.

20. Click **Next**.

21. Click **Submit**.

When provision starts, the vFiles will appear in the Datasets panel. Select the vFile and navigate to the Status tab to see the progress of the job. When provisioning is complete, you can see more information on the Configuration tab.

Related Topics

- [Unstructured Files Environment Requirements](#)
- [Linking Unstructured Files](#)

Managing vFiles

Overview

This article is used to cover steps on adding an additional mount to an existing vFile and outlining post provision/migration ownership rules.

Adding an Additional Mount

1. Login to the **Delphix Management** application.
2. Select the **vFile** you want to edit.
3. From the Actions menu (...) select **Disable**.
4. In the **Configuration** tab select the **Source** sub-tab, a Pencil icon will appear on the right of the Additional Mount Points.
5. Click on the **Pencil** to edit the mount points.
6. In the **Additional Mount Points** window select the **Plus** icon to add additional mount points.
7. Click the **checkmark** to save.
8. From the Actions menu (...) **Enable** the vFile.

Post Provision/Migration Ownership Rules

When a new vFile VDB is provisioned, the ownership is changed to match the Environment User anytime a VDB is enabled. This could cause conflict in the ownership of existing files in a case where the VDB has just been migrated to a new host – the VDB files on the new host will now be owned by the new Environment User.

Related Topics

- [Unstructured Files and App Data](#)
- [Linking Unstructured Files](#)

vFiles Best Practices and Common Pitfalls

Overview

This document is the implementation guide for the best practices of implementing data source integrations using the Delphix vFile functionality provided through the AppData toolkit. Since "unstructured files", or vFiles, typically is implemented for non-DBMS file types, this implementation will require a certain degree of configuration and scripting in order to function. This means that while you do get some of the main functionality of the Delphix Engine, there are also quite a few pitfalls that may deter customers from wanting this type of configuration. Using vFiles and scripts to implement data source ingestion is a workaround that should only be used if you and the customer are fully aware of the best practices and limitations as called out below.

Best Practices and Implementation

In order to be successful with this workaround, please keep in mind these best practices which will enable you to be successful in implementing a vFile configuration for data sources.

- Write a script to put the RDBMS into 'Backup' mode to prepare it for the next set of scripts Delphix will perform
- Create tablespaces in the RDBMS which dxtoolkit will take snapshots of while the database is in backup mode
- After any vFile provisioning job, script a recovery of the RDBMS instance

Common Use Cases

These use cases are the typical scenarios for which you would use this workaround.

Unsupported DBMS Configuration

This functionality is most appropriate when a customer is using an unsupported data source and wants a quick MVP to demonstrate the value of Delphix. It is critical to understand the pitfalls of this implementation and what value propositions that this excludes so that the customer is aware of what capabilities our core integrations provide.

Common Pitfalls

These pitfalls of functionality should be considered whenever implementing vFile scripted ingestion. Refer to the table for a quick reference of pitfalls.

Pitfall	Description
LogSync	Not Available
Point in Time Recovery	Not Available
Clustering	Not Available
User Interface	Not Available

Pitfall	Description
Production Downtime	Yes
Use of dxtoolkit	Requires Training

LogSync and Point in Time Provisioning

Since these scripts are taking full backups of the database, there will not be logs available for us to sync with. Therefore, point in time provisioning is not available to any vFile implementation.

Manual Operations

Currently, customers using this workaround have to perform the scripts manually. This creates operational overhead for customers, particularly their DBA team.

Clustering

Using vFiles is not available for clustered instances, and there is currently no workaround for this type of configuration.

Lack of Graphic User Interface (GUI)

Scripts perform all the core functionality that is typically carried out in the Delphix GUI. This implementation precludes the usage of the main GUI and is only accessible via the Command Line Interface (CLI).

Production Downtime

As the database enters backup mode during the ingestion process, this workaround causes the production source to go down while Delphix ingests their data.

Understanding and Implementing Scripts and Dxtoolkit

Implementing these scripts requires an understanding of dxtoolkit and the data source which you are trying to integrate with. Typically, we estimate the learning period of this to be about 1 - 2 months to fully understand how to use these tools and how to best implement the scripts for any given source.

Conclusion

Using vFile functionality with scripts is a quick way to demonstrate the value of the Delphix Engine easily. With a little scripting and Delphix knowledge, you can engineer ingestion with dxtoolkit to take snapshots of the entire database. However, there are several pitfalls that every customer should know about and understand before moving forward with this type of implementation. If you have any questions, please reach out to Product Management.

Related Topics

- [Unstructured Files and App Data](#)

Delphix Engine Plugin Management

Plugin Management

This section provides information on the data management plugin deployed on the Delphix Engine. Users can directly upload their Plugins to the Delphix Management application with a single click, thus eliminating the process of uploading the Plugin using the command. Plugin upload via the Delphix Management application helps users to build their Plugin library in a simple and easy-to-use manner.

The Plugin screen allows users to view a structured or formatted view of their selected Plugin and it provides a single click option for users to delete a Plugin.

 Only Admin users are able to upload or delete Plugins. Standard users are only able to view Plugin information.

Plugin Types

There are two types of plugins the Delphix engine supports Lua/Python plugins and platform plugins. Because support for these types of plugins was written at different times, the upgrade and replication workflows will differ depending on what the plugin type is. To figure out what type of plugin this is, follow the instructions below:

1. Login to the **Delphix Management** application.
2. Select **Manage > Plugins**.
3. Select the plugin you want to check and look for the type field in the details section on the right.
4. If the type is Toolkit this means the plugin is a Lua plugin. If the type is Plugin, then the plugin is either a platform plugin or a python plugin.

Installation of Plugin

Plugins can be uploaded onto the Delphix Engine only by an admin.

Prerequisites


- Make sure you have access to a copy of the plugin from the computer you will be installing.
- Verify with the plugin author that the version of the plugin you have is compatible with your Delphix Engine version.

Procedure

To add or upload a new Plugin complete the instructions below:

1. Login to the **Delphix Management** application.
2. Select **Manage > Plugins**.

3. To add or upload a Plugin, click the  icon. This opens the Upload or Upgrade Plugin dialog.



4. To start using the plugin, refresh target and source environments to discover any repositories.

Upgrading a Plugin

When a new version of a plugin is released, the new plugin can only be uploaded onto the Delphix Engine by an Admin.

Prerequisites

- Make sure you have access to a copy of the plugin from the computer you will be installing.
- Verify with the plugin author that the new version of the plugin is compatible with your current version, and with your Delphix Engine version.
- If the already-installed plugin does not use the Lua language, then you don't have to disable any dSources or VDBs. However, be aware that you will not be able to perform any Delphix operations on them while the plugin is being upgraded.

Procedure

To upgrade a Plugin just follow the instructions above on how to upload a Plugin. The newly-uploaded plugin will be used to upgrade any already-present version(s) of the same plugin.

Notes on Replication with Objects Administered by a Plugin

Replication works the same way for objects created with plugins and objects created using natively supported data types. But, complications and problems can arise when a replicated plugin's version does not match the version of the installed plugin.

Replication Recommendation

As discussed in the section on [Delphix Replication Overview](#), we highly recommend that replication target engines be used solely to hold replicated objects, and that failover is only done in disaster recovery situations. Or, at the least, make sure that the versions of replicated plugins match the versions of plugins that are already installed on the replication target engine.

There are specifically two operations done after replication that will be listed below.

Provisioning from Replicated Data Sources or VDBs

Replica provisioning is only guaranteed to work when the source object's plugin is already also installed on the target engine with the identical version.

Installed Plugin Version	Result
None	Replica provisioning will fail.
Lower than replicated version	Replica provisioning will fail.
Same as replicated version	Replica provisioning will succeed.
Higher than replicated version	<p>If the replicated plugin uses the Lua language, replica provisioning will fail.</p> <p>If the replicated plugin does not use the Lua language, replica provisioning may or may not succeed, depending on whether the installed plugin is compatible with the replicated version.</p>

More generic information on how to perform a replica provision can be found [here](#).

Failing Over a Replica

As mentioned above, we strongly recommend that replication target engines do not have any non-replica objects. Although the replica failover process will always work, the failed-over objects may or may not be fully functional if there are non-replica plugin-administered objects already on the same engine. In some cases, you may have to delete some objects, and their associated plugins.

Installed Plugin Version	Result
None	The replicated plugin and its objects will be moved into the live namespace. All failed-over objects will be fully functional.
Lower than replicated version	The failed-over objects may become "inactive", and any Delphix operations run on them will fail. In this case, you can try to upgrade the existing plugin so that the version matches the replicated inactive plugin. This may or may not work, depending on whether there are any conflicts between any objects.

Installed Plugin Version	Result
Same as replicated version	The failed-over objects will be fully functional.
Higher than replicated version	The failed-over objects will become "inactive", and any Delphix operations run on them will fail. In this case, you can try to upgrade the inactive plugin (see instructions below). Again, this may or may not work, depending on whether there are conflicts between any objects, and on whether the installed plugin knows how to interact with the inactive plugin.

More generic information on how to perform an actual failover can be found [here](#).

Upgrading an Inactive Plugin

If a failed-over plugin becomes "inactive" (see details above), you may be able to activate the plugin by upgrading it to match the version of the installed plugin.

Procedure

To upgrade an inactive Plugin complete the instructions below:

1. Login to the **Delphix Management** application.
2. Select **Manage > Plugins**.
3. In the drop-down menu on the left, select the plugin that is inactive.

Plugins

+
↕
🗑️

Pretty Name	Version	Namespace
Unstructured Files	1.0.0	
nix_staged_python	2.1.0	
nix_staged_python	2.0.0	

```

▼ object {23}
  externalVersion : 2.0.0
  buildNumber : 2
  ▶ virtualSourceDefinition {2}
  ▶ linkedSourceDefinition {2}
  ▶ discoveryDefinition {8}
  ▶ upgradeDefinition {1}
    entryPoint : operations.nix_staged:staged
    sourceCode : *****
  ▶ manifest {21}
    status : INACTIVE

```

4. Click the upgrade button  which should be clickable now to activate the inactive plugins. Make sure to confirm the upgrade by selecting OK.



5. After the upgrade is successful the inactive plugin of the lower version will not exist in the plugin list.

 If the inactive plugin is a higher version, the active plugin needs to be upgraded using the normal upgrade procedure instead.

Related Topics

- [Db2 Plugin Installation and Upgrade](#)

Unstructured Files Hook Operation Notes

Shell Operations

RunCommand Operation

The RunCommand operation runs a shell command on a Unix environment using whatever binary is available at `/bin/sh`. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command line interface (CLI) to aid in debugging.

If successful, the shell command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Examples of RunCommand Operations

You can input the full command contents into the RunCommand operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"

if test -d "$remove_dir"; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

If a script already exists on the remote environment and is executable by the environment user, the RunCommand operation can execute this script directly.

```
/opt/app/oracle/product/10.2.0.5/db_1/dbs/myscript.sh "$ARG_ENVIRONMENT_VARIABLE"
"second argument in double quotes" 'third argument in single quotes'
```

RunBash Operation

The RunBash operation runs a Bash command on a Unix environment using a `bash` binary provided by the Delphix Engine. The environment user runs this Bash command from their home directory. The Delphix Engine captures and logs all output from this command. If the script fails, the output is displayed in the Delphix Management application and command line interface (CLI) to aid in debugging.

If successful, the Bash command must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of RunBash Operations

You can input the full command contents into the RunBash operation.

```
remove_dir="$DIRECTORY_TO_REMOVE_ENVIRONMENT_VARIABLE"
```

```
# Bashisms are safe here!
if [[ -d "$remove_dir" ]]; then
    rm -rf "$remove_dir" || exit 1
fi

exit 0
```

Shell Operation Tips

Using `nohup`

You can use the `nohup` command and process backgrounding from resource in order to "detach" a process from the Delphix Engine. However, if you use `nohup` and process backgrounding, you MUST redirect `stdout` and `stderr`.

Unless you explicitly tell the shell to redirect `stdout` and `stderr` in your command or script, the Delphix Engine will keep its connection to the remote environment open while the process is writing to either `stdout` or `stderr`. Redirection ensures that the Delphix Engine will see no more output and thus not block waiting for the process to finish.

For example, imagine having your `RunCommand` operation background a long-running Python process. Below are the bad and good ways to do this.

Bad Examples

- `nohup python file.py & # no redirection`
- `nohup python file.py 2>&1 & # stdout is not redirected`
- `nohup python file.py 1>/dev/null & # stderr is not redirected`
- `nohup python file.py 2>/dev/null & # stdout is not redirected`

Good Examples

- `nohup python file.py 1>/dev/null 2>&1 & # both stdout and stderr redirected, Delphix Engine will not block`

Other Operations

RunExpect Operation

The `RunExpect` operation executes an Expect script on a Unix environment. The Expect utility provides a scripting language that makes it easy to automate interactions with programs which normally can only be used interactively, such as `ssh`. The Delphix Engine includes a platform-independent implementation of a subset of the full Expect functionality.

The script is run on the remote environment as the environment user from their home directory. The Delphix Engine captures and logs all output of the script. If the operation fails, the output is displayed in the Delphix Management application and CLI to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunExpect Operation

Start an `ssh` session while interactively providing the user's password.

```
spawn ssh user@delphix.com
expect {
  -re {Password: } {
    send "${env(PASSWORD_ENVIRONMENT_VARIABLE)}\n"
  }
  timeout {
    puts "Timed out waiting for password prompt."
    exit 1
  }
}
exit 0
```

RunPowershell Operation

The RunPowershell operation executes a PowerShell script on a Windows environment. The environment user runs this shell command from their home directory. The Delphix Engine captures and logs all output of the script. If the script fails, the output is displayed in the Delphix Management application and command line interface (CLI) to aid in debugging.

If successful, the script must exit with an exit code of `0`. All other exit codes will be treated as an operation failure.

Example of a RunPowershell Operation

You can input the full command contents into the RunPowershell operation.

```
$removedir = $Env:DIRECTORY_TO_REMOVE

if ((Test-Path $removedir) -And (Get-Item $removedir) -is [System.IO.DirectoryInfo])
{
  Remove-Item -Recurse -Force $removedir
} else {
  exit 1
}
exit 0
```

Unstructured Files Environment Variables

Operations that run user-provided scripts have access to environment variables. For operations associated with specific dSources or virtual databases (VDBs), the Delphix Engine will always set environment variables so that the user-provided operations can use them to access the dSource or VDB.

dSource Environment Variables

Environment Variable	Description
DLPX_DATA_DIRECTORY	Path where linked-staged database is mounted

VDB Environment Variables

Environment Variable	Description
DLPX_DATA_DIRECTORY	Path where virtual database is mounted

Virtualizing Oracle E-Business Suite

This section covers the following topics:

- [Oracle E-Business Suite and Delphix Conceptual Overview](#)
- [Oracle E-Business Suite Support Matrix](#)
- [EBS Plugin Installation](#)
- [Upgrade EBS Objects to 6.0](#)
- [Upgrade Objects to 5.3](#)
- [Upgrade EBS objects to 5.2](#)
- [Oracle EBS R12.2](#)
- [Managing Data Operations for Virtual EBS Instances](#)
- [Virtual EBS Instance Recipes](#)

Oracle E-Business Suite and Delphix Conceptual Overview

This topic provides a conceptual overview of the integration between the Delphix Engine and Oracle E-Business Suite (EBS).

Why Use the Delphix Engine with E-Business Suite?

Oracle E-Business Suite offers a multitude of enterprise resource planning (ERP), supply chain management (SCM), and customer relationship management (CRM) applications useful for running a business. Typical EBS instances are consumed horizontally across an organization but managed on-site by an IT organization. Most businesses hire consultants or in-house developers to customize these instances to fit business needs.

Production EBS instances are critical to many business processes; EBS is expected to be always available and functioning correctly. An IT organization's EBS team is responsible for maintaining this production instance: this team monitors the instance and manages how software modifications are integrated. Common EBS software modifications include patches provided by Oracle or new application logic provided by consultants or in-house developers.

However, before making a software modification to their production EBS instance, most EBS teams vet the modification. These teams typically maintain a series of non-production copies of their production EBS instance to use for the vetting process. To vet a modification, the EBS team will modify a non-production instance in an attempt to identify bugs in the modification that might lead to incorrect functionality, performance regressions, or downtime. Once a modification has been installed and tested successfully, it is considered vetted and the EBS team make the same modification to the production instance.

The problem with this process is that maintaining non-production copies of EBS is costly. EBS teams not only need to provision and maintain resources for each copy, they also must update each copy to match production every time production itself is modified. This process takes two hours at a minimum and can take up to a week or more.

The Delphix Engine can alleviate much of the pain around managing non-production copies of Oracle E-Business Suite. You can link a production EBS instance to the Delphix Engine and use the engine to provision virtual copies. These virtual copies lower storage costs associated with non-production. They also drastically reduce the amount of active attention required from the EBS team. When a change is made to the production EBS instance, the Delphix Engine can synchronize its copy; all virtual EBS instances can then be refreshed to match this new version of the production. The process of refreshing a virtual EBS instance is fully automated, so an EBS team no longer needs to invest copious amounts of time refreshing non-production instances manually.

In addition, the Delphix Engine can take snapshots of virtual EBS instances in known healthy states so that you can easily roll back risky modifications while vetting them. This functionality provides modern version control semantics for a process that previously had none.

Overview of Linking and Provisioning EBS

Oracle E-Business Suite is primarily comprised of a database service and a plethora of application services. The Delphix Engine supports the linking and provisioning of all EBS datasets, including the database technology stack (dbTechStack), database, and application files (appsTier). Note that the Delphix Engine can also manage custom extensions and plug-ins.

The process of linking EBS data involves creating multiple dSources:

- a dSource for the dbTechStack
- a dSource for the Oracle database used by EBS
- a dSource for the appsTier

These dSources are collectively referred to as EBS dSources. They are also the source datasets from which you can provision virtual EBS instances.

The process of provisioning a virtual EBS instance involves provisioning from each of these dSources separately to the proper environments in sequence. You can add custom configuration logic to the Delphix Engine for each EBS instance such that the linking, provisioning, and refresh processes are fully automated.

 The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

dbTechStack Dataset


The source dbTechStack is linked using the Delphix Engine's EBS support: the linking process automatically runs pre-clone logic to ensure EBS configuration is always appropriately staged at the time of data capture. When you provision EBS, the Delphix Engine automates post-clone configuration such that a copy of the dbTechStack is available for use on the target dbTier server with no additional effort. You can add this copy of the dbTechStack to the Delphix Engine as an Oracle installation home and use it to host an EBS virtual database (VDB).

Database Dataset

The database dSource is linked using the Delphix Engine's support for Oracle databases. This dSource contains database data files that EBS is currently using. For more information about managing Oracle databases. When you provision EBS, you will use the Delphix Engine to set up a copy of the EBS database on the target dbTier server. This copy of the database will be used to back virtual EBS instance's appsTier.

appsTier Dataset

The appsTier is linked using the Delphix Engine's EBS support: the linking process automatically runs pre-clone logic to ensure EBS configuration is always appropriately staged at the time of data capture. When you provision EBS, the Delphix Engine will automate post-clone configuration such that a copy of the appsTier is available for use on the target appsTier server. This virtual copy of the appsTier will connect to the provisioned EBS virtual database (VDB).



How Does It Work?

The Delphix Engine's linking and provisioning logic follows the Oracle-recommended cloning procedures outlined in the following Oracle Support documents found at <http://docs.oracle.com>:

- Cloning Oracle E-Business Suite Release 12.2 with Rapid Clone (Doc ID 1383621.1)
- Cloning Oracle Applications Release 12 with Rapid Clone (Doc ID 406982.1)

Related Topics

- [Virtualizing Oracle E-Business Suite](#)

Oracle E-Business Suite Support Matrix

Important Notes

- Delphix Support Policies specifically list Major and Minor release coverage. If a minor release is listed as covered, then all patch releases under that minor release are certified.
- Support applies to corresponding versions of Community Enterprise Operating System (CentOS) / Oracle Linux (OL)
- 64-bit OS support only

 For specific plugin version release notes, see [EBS 12.2 Plugin Release Notes](#).

Key:

Color	Supported?
	Yes
	No
	Not Applicable

Red Hat Enterprise Linux (RHEL)

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 5.5		
RHEL 5.6		
RHEL 6.4		
RHEL 6.5		
RHEL 6.6		

Supported OS Version	Supported EBS 12.2 Apps version	
	Oracle 11g/12c	Oracle 19c
RHEL 6.7	✓	N/A
RHEL 6.8	✓	N/A
RHEL 6.9	✓	N/A
RHEL 7.0	✓	N/A
RHEL 7.1	✓	N/A
RHEL 7.2	✓	N/A
RHEL 7.3	✓	N/A
RHEL 7.4	✓	N/A
RHEL 7.5	✓	✓
RHEL 7.6	✓	✓
RHEL 7.7	✓	✓
RHEL 7.8	✓	✓
RHEL 7.9	✓	✓
RHEL 8.0	✗	✓
RHEL 8.1	✗	✓
RHEL 8.2	✗	✓

	Supported EBS 12.2 Apps version	
Supported OS Version	Oracle 11g/12c	Oracle 19c
RHEL 8.3	✗	✓
RHEL 8.4	✗	✓
RHEL 8.5	✗	✓
RHEL 8.6	✗	✓

Reference
 Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)

Solaris SPARC (x86 and SPARC)

	Supported EBS 12.2 Apps version
Supported OS Version	Oracle 12.1.0.2
Solaris 11 U4	✓

Plugin/Delphix Engine Compatibility

	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
Supported Delphix Engine	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.0.0	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.1.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗

Supported Delphix Engine	Plugin Version for EBS 12.2 with Oracle 11g/12c										Plugin Version for EBS 12.2 with Oracle 19c										
	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	1.8.2	1.8.3	1.8.4	2.0.0	2.0.1	3.0.0	3.0.1	3.0.2	3.0.3	4.0.0	
6.0.2.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.3.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.4.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.5.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.6.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.7.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.8.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
6.0.9.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✗	✗	✗	✗	✗
6.0.10.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✗	✗	✗	✗
6.0.11.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗	✗
6.0.12.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✗	✗
6.0.13.0	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓

Plugin Upgrade Path for EBS 12.2 Customers

i The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

New plugin installation for EBS customers

Upload 3.0.0/3.0.1/3.0.2/3.0.3/4.0.0 plugin for EBS 12.2 customers.

Existing EBS customers with different plugin version

Your Version	Recommended Upgrade Path to 3.0.3
1.8.2	Upgrade to 2.0.0 or 2.0.1 and follow the upgrade path below.
1.8.3	
1.8.4	
2.0.0	Upgrade to 3.0.0, 3.0.1, 3.0.2, or 3.0.3 the latest plugin version.
2.0.1	
3.0.0	Upgrade directly to the latest 4.0.0 plugin version.
3.0.1	
3.0.2	
3.0.3	

For example, upgrade path for EBS 12.2 customer with existing plugin version 1.8.2 who wants to upgrade to 3.0.0 plugin version : 1.8.2 or 1.8.3 or 1.8.4 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2, or 3.0.3 → 4.0.0.

- Customers should skip a minor plugin 1.8.2/1.8.3/1.8.4 version upgrade.
- If EBS 12.2 customers are on plugin version 1.8.2, the plugin upgrade path will be : 1.8.2 → 2.0.0 or 2.0.1 → 3.0.0, 3.0.1, 3.0.2 or 3.0.3 → 4.0.0.

Unsupported EBS Versions and Features

- Support SSL enabled Environment
- V2P

Related Topics

- [Virtualizing Oracle E-Business Suite](#)
- [EBS 12.2 Plugin Release Notes](#)

EBS Plugin Installation

The EBS Plugin is provided as a zip file on the Delphix download site. Follow the steps below to download the plugin:

1. In the web browser, go to the Delphix [download](#) site.
2. Login to the download site using email and password credentials.
3. Navigate to the **Delphix Product Releases** folder and select the required Delphix Engine version.
4. Go to the `Plugins> EBS> EBS12.2> EBS_Toolkit_<version_number>_for_EBS12.2.zip` folder, then download the zip file, unzip it, and copy both the .json files at your preferred location.

 For a particular version of the EBS plugin, there are two .json files that will be downloaded and the same would be uploaded while installing or upgrading the plugin. Also, note that you need to upload these .json files in sequential order. For example, for EBS12.2 2.0.1 plugin version you need to upload *ebs122-db2.0.1.json* first and then *ebs122-app2.0.1.json*.

 The plugin versions for AppsTier and DB Tech Stack need to be the same, at installation and through the upgrade. Installation and usage of different versions of these plugins are not supported and may encounter operational issues.

For plugin upgrade and installation steps, refer to [Delphix Engine Plugin Management](#). Also, refer to [EBS Support Matrix](#) for upgrade paths.

Related Links

- [EBS Support Matrix](#)
- [Delphix Engine Plugin Management](#)

Upgrade EBS Objects to 6.0

This section focuses on the steps to be taken post-upgrade to Delphix 6.0 release for both EBS 12.1 and EBS 12.2 for all dSources and VDBs present on the Delphix Engine.

Uploading a Plugin

After upgrading to the required DE version 6.0.3 or later, if you need to upgrade or install a new plugin, follow the steps mentioned in [EBS Plugin Installation](#).

Post Upgrading to 6.0.3.0 or later

EBS 12.2 plugin

This section focuses on the steps taken after upgrading to the Delphix 6.0.3.0 release or later, for all the users who intend to use plugin version 2.0.0 or 2.0.1 for EBS 12.2.

Pre-requisites for the existing users:

To successfully upgrade to Delphix Engine 6.0.3.0 or later with EBS 12.2 plugin version 3.0.0, the following conditions must be fulfilled:

- The user must be on or upgrade to the 2.0.0/2.0.1 version before moving to 3.0.0 on the Delphix Engine. Refer [Plugin Upgrade Path for EBS 12.2 Customers](#) for all the recommended upgrade paths.
- All datasets on Delphix Engine must be in a disabled state.

Post Upgrade Steps:

1. Verify the plugin version under **Manage > Plugins** section. It should show as 3.0.0.
2. Refresh all the environments on the Delphix Engine.
3. Users who intend to use low-privileged user, (for example, "delphix_os") should follow the environment configuration mentioned [here](#) and steps 4 and 5 below.
4. All the EBS dSources and VDB/vFiles should edit the Environment User under the Configuration → Source to "delphix_os".
5. Users who intend to use low-privileged user, must update the "Privileged OS User" field on the dSource/VDB provision wizard to high privileged user (user who should own the file system).
6. Enable all the EBS dSources - dbTechStack, Database, and appsTier.
7. Take new snapshots for all the enabled EBS dSources.
8. Enable all the EBS VDB/vFiles as per the following sequence: (1) dbTechStack (2) Database (3) appsTier
9. Under **Custom > Configuration** tab for dbTechStack vFiles, update the value for the "Target APPS Password" field to the value similar to the "Target APPS Password" field as provided under appsTier vFiles.
10. Modify the database hooks.
Follow the steps under **Provisioning the Oracle Database** on [Provisioning a Virtual EBS R12.2 Instance](#) to configure the hooks.
11. Take new snapshots for all the enabled EBS VDB's.
12. To refresh the VDB's, follow the steps and instructions on [Refreshing a Virtual EBS Instance](#).

Related Topics

- [Virtualizing Oracle E-Business Suite](#)

Upgrade Objects to 5.3

Post Upgrade to 5.3.2.0

This section focuses on steps to be taken post-upgrade to the Delphix 5.3.2.0 release for all customers who intend to use WebLogic password change feature for EBS 12.2 release.

1. Disable all the virtual datasets created on Delphix Engine in the following order:
 - a. appsTier vFiles
 - b. VDB utilized by your EBS instance
 - c. dbTechStack vFiles
2. Disable all the EBS dSource created on Delphix Engine once you have disabled all three EBS virtual datasets in the above step.
3. Upgrade Delphix Engine to 5.3.2.0 release and upload EBS 1.4.0 plugin release.
Go to Manage → Plugins
 - Upload a plugin file to upgrade existing toolkit by clicking plus (+) symbol.
4. Refresh all host environments added on the Delphix Engine
Go to Manage → Environments
 - From the symbol (...) located in the upper left-hand corner, select Refresh All.
5. Change the Weblogic admin password on EBS appsTier vFile.
 - Select the appsTier vFiles hosting your virtual EBS database.
 Go to Configurations → Custom
 - Update “**Target Weblogic AdminServer Password**” field with the password you would like to set for EBS appsTier.
6. Enable all the EBS dSource created on Delphix Engine.
7. Enable all the virtual datasets created on Delphix Engine in the following order:
 - a. dbTechStack vFiles
 - b. VDB utilized by your EBS instance
 - c. appsTier vFiles
8. Verify the changed WebLogic password by logging into the WebLogic console.

Post Upgrade to 5.3.1.0

This section focuses on steps to be taken post-upgrade to the Delphix 5.3.1.0 release for all customers who intend to use sudo privileges for EBS 12.1 release.

1. Create a low-privileged OS account (delphix_os) in source and target environments.
2. Steps and configuration required for a low-privileged account to perform delphix privileged operations are outlined in [Requirements for Unix Environments](#).
3. Push the privilege elevation script "DLPX_DB_EXEC" to both and target environments as per the steps outlined in [Privilege Script](#).
4. For all 12.1 EBS appsTier Dsource :

Go to Configuratio>Source

- Update Environment User value to low privileged account (delphix_os).

Go to Configuration>Custom

- Update Privileged OS Account with the value for high privileged user (applmgr).

5) For all 12.1 EBS appsTier vFile :

Go to Configuration>Source

- Update Environment User value to low privileged account (delphix_os).

Go to Configuration>Custom

- Update Privileged OS Account with the value for high privileged user (aplmgr).

6) For both source and target environment:

Go to Manage>Environments

- Add low privileged OS user (delphix_os) in Environments Users section by clicking plus (+) symbol.
- Mark low-privileged OS user (delphix_os) as the primary user.
- Remove the high privileged OS user (aplmgr) from the Environments Users section.

Related Topics

- [Virtualizing Oracle E-Business Suite](#)

Upgrade EBS objects to 5.2

This section focuses on steps to be taken post-upgrade to the Delphix 5.2 release for all dSources and VDBs present on the Delphix Engine.

1) For all 12.1 EBS apps Tier Dsource :

Go to Configuration → Custom

- Update context_name with the Context value of the application tier
- Update apps password with the value of the dsource.
- Take a new snapshot after updating the above parameters

2) For all 12.1 and 12.2 EBS Database VDB:

- Update the pre snapshot hook as follows

```
# NOTE: Ensure the below environment variables will be set up correctly by the shell.
If not, hardcode or generate the values below.
# CONTEXT_NAME=${ORACLE_SID}_${hostname -s}
# APPS_PASSWD=<source apps passwd>
# TARGET_APPS_PASSWD=<new apps password>

timeout=3600
waittime=0
. ${ORACLE_HOME}/${CONTEXT_NAME}.env

testAppsPassword() {
    local passwordInQuestion=$1
    ERROR=`sqlplus "apps/${passwordInQuestion}" <<< "exit;"`

    grep ORA-01017 <<< ${ERROR} >/dev/null && return 1
    return 0
}

testAppsPassword ${SOURCE_APPS_PASSWD}
testResult=$?

if [[ ${testResult} == 0 ]]; then
    APPS_PASSWD=${SOURCE_APPS_PASSWD}
else
    APPS_PASSWD=${TARGET_APPS_PASSWD}
fi

while [[ -f "${ORACLE_HOME}/.delphix_adpreclone.lck" || "$(ps -Ao args | grep "^${ORACLE_HOME}" | grep -v grep | grep "adpreclone")" ]]; do
if [[ $waittime -gt $timeout ]]; then
    echo "Another adpreclone process is still running from last 60 mins.Delphix cannot proceed until it is complete. Exiting Now."
    exit 1
fi
```

```
    echo " Another adpreclone process is running. waiting for the process to
complete before starting adpreclone of the database...."
    (( timeout += 10 ))
    sleep 10
done
echo "No other adpreclone process is running on database, proceeding ...."
. ${ORACLE_HOME}/${CONTEXT_NAME}.env
${ORACLE_HOME}/perl/bin/perl ${ORACLE_HOME}/appsutil/scripts/${CONTEXT_NAME}/
adpreclone.pl database <<-EOF
${APPS_PASSWD}
EOF
```

3) For all 12.1 EBS apps Tier VDB :

Go to Configuration → Custom

- Update system password with the system password corresponding to the database.
- Take a new snapshot after updating the above parameters

Related Topics

- [Virtualizing Oracle E-Business Suite](#)

Oracle EBS R12.2

This section covers the following topics:

- [Source EBS R12.2 Instance Requirements](#)
- [Virtual EBS R12.2 Instance Requirements](#)
- [Remote Logging](#)
- [Preparing a Source EBS R12.2 Instance for Linking](#)
- [Linking a Source EBS R12.2 Instance](#)
- [Preparing Target EBS R12.2 Environments for Provisioning](#)
- [Provisioning a Virtual EBS R12.2 Instance](#)

Source EBS R12.2 Instance Requirements

The Delphix Engine supports linking a variety of versions and configurations of Oracle E-Business Suite R12.2. Below, detailed compatibility notes are outlined.

Note that minor releases of EBS are not certified for Delphix Engine compatibility individually; major release support implies support for any minor release.

To review the requirements for provisioning a virtual EBS R12.2 instance from the Delphix Engine, see [Virtual EBS R12.2 Instance Requirements](#).

Operating Systems

- Linux – SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris - SPARC 9, 10, and 11

AIX and HP-UX Not Supported

The Delphix Engine does not support linking EBS instances running on AIX or HP-UX.

dbTier Requirements

Supported Topologies

- Oracle single-instance (SI) dbTechStack and Database
 - Oracle11gR2 v11.2.0.4
 - Oracle12cR1 v12.1.0.2
 - Oracle19c MT v19.3.0.0
- Oracle RAC dbTechStack and Database
 - Oracle11gR2 v11.2.0.4
 - Oracle12cR1 v12.1.0.2

Caveats

The Delphix Engine does not provide support for linking an EBS R12.2 instance utilizing custom context variables maintained in the EBS context file.

Related Topics

- [Virtual EBS R12.2 Instance Requirements](#)
- [Oracle EBS R12.2](#)

Virtual EBS R12.2 Instance Requirements

The Delphix Engine supports provisioning a variety of versions and configurations of Oracle E-Business Suite R12.2. Below, detailed compatibility notes are outlined.

Note that minor releases of EBS are not certified for Delphix Engine compatibility individually: major release support implies support for any minor release.

To review the requirements for linking an EBS R12.2 instance to the Delphix Engine, see [Source EBS R12.2 Instance Requirements](#).

Operating Systems

- Linux - SLES 9, 10, 11; RHEL 5, 6, 7, OL 5, 6
- Solaris - SPARC 9, 10, and 11

AIX and HP-UX Not Supported

The Delphix Engine does not support provisioning EBS instances to AIX or HP-UX.

dbTier Requirements

Supported Topologies

- Oracle SI dbTechStack and Database
 - Oracle11gR2 v11.2.0.4
 - Oracle12cR1 v12.1.0.2
 - Oracle19c MT v19.3.0.0

Oracle RAC Not Supported

The Delphix Engine does not support provisioning Oracle RAC dbTechStack for use with Oracle E-Business Suite.

However, you may provision an Oracle SI dbTier from a linked EBS RAC dbTier instance. During the provisioning process, the Delphix Engine will relink the dbTechStack for use with an Oracle SI database and scale down the Oracle RAC database to an Oracle SI database.

appsTier Requirements

Supported Topologies

- Single-node appsTier
- Multi-node appsTier with a shared APPL_TOP

Non-shared APPL_TOP Not Supported

The Delphix Engine does not provide support for provisioning a multi-node appsTier where the APPL_TOP is not shared between nodes.

appsTier Topology Is Configurable

The appsTier topology of the virtual EBS instance does NOT need to match the appsTier topology of the source EBS instance. The Delphix Engine will automate scale up or scale down logic as required during provisioning.

Caveats

The Delphix Engine will only register the default WLS-managed server for each server type during multi-node appsTier provisioning. Additional WLS managed servers need to be registered manually or as a **Configure Clone** hook.

Related Topics

- [Source EBS R12.2 Instance Requirements](#)
- [Oracle EBS R12.2](#)

Remote Logging

Remote Logging for EBS 12.2 Plugin adds the following options to provide enhanced logging:

- set debug log levels
- set active log file size
- set log retention levels

You can modify the parameters defined in the `ebs_plugin.conf` file located in remote target host at `<Toolkit directory>/<Delphix_COMMON>/plugin/ebs122-vsdk-db_9f7f5a63-1da7-495c-80df-e0fa61941342/ebs_plugin.conf` for DbTechStack and `<Toolkit directory>/<Delphix_COMMON>/plugin/ebs122-vsdk-app_403e7034-0028-4200-ad0b-15031f0f8871/ebs_plugin.conf` for appsTier to configure the logging mechanism.

Setting Log Levels (Level)

You can set log levels for plugin-generated logs by setting the parameter `level` in the `ebs_plugin.conf` file. Higher logging levels will help to expedite debugging issues. There are two levels of logging: info and debug.

Log Level Description

- `level=INFO`: This level will print only informational logs. This is the default log level.
- `level=DEBUG`: This level will print informational logs and debug statements.

Setting Maximum Size of Active Log File (logFileSize)

You can set this parameter inside the `ebs_plugin.conf` file to set the maximum size of the active log file in MB. Once this limit is reached, the plugin rotates the log as per the retention property defined below. The minimum value of this parameter is 1 MB and the maximum value is 10 MB. This parameter only takes in a positive integer value. The default value is 1 MB.

Setting Log Retention Levels (Retention)

You can set a retention level for diag logs (`diag.log`) using parameter `retention` in the `ebs_plugin.conf` file. As per this parameter, the log files are moved (archived), renamed, or deleted once they reach the value set in the `logFileSize` parameter. New incoming log data is directed into a new fresh file (at the same location).

By default, this value will be set to a minimum value of 2. The user can change this value and set its value within the range 2 and 50. For example; if retention is set to 4, the plugin will have the following log files: `diag.log`, `diag.log.1`, `diag.log.2`, `diag.log.3`, `diag.log.4`.

- File `diag.log` is the active log file.
- File `diag.log.1` is the most recent archive log file
- File `diag.log.4` is the oldest one

ebs_plugin.conf file

```
#
# This flag will set the debug level of plugin logs on the remote server. There are
two levels:
```

```
# Info
# Debug
# The above are the only valid values that can be assigned. If any other value is
# assigned, the plugin will set the
# default level as Info. the string is case insensitive so info, Info and INFO are
# acceptable.
#
level=INFO

#
# This parameter will set the maximum size of the active log file in MB. Once this
# limit is reached, the plugin will
# rotate the log as per the retention property defined below. The minimum value of
# this parameter is 1 MB and maximum
# value is 10 MB. This parameter only takes in a positive integer value. The default
# value is 1 MB
# logFileSize=<positive integer>
#
logFileSize=1

#
# Whenever the size of diag.log file exceeds the value provided by logFileSize
# parameter then the plugin will
# rename the active log file to diag.log.<number> and a new log file with name
# diag.log will be
# generated. For example, if LogRetention is set to 4, the plugin will have the
# following log files: diag.log,
# diag.log.1, diag.log.2, diag.log.3, diag.log.4.
# File diag.log.4 will be the oldest one.
# File diag.log.1 will be the most recent archive log file.
# File diag.log will be the active log file.
# The minimum and default value for retention flag is 2.
# retention=<positive integer>
#
retention=2
```

Preparing a Source EBS R12.2 Instance for Linking

This topic outlines the prerequisites for linking an EBS R12.2 instance to the Delphix Engine.

Ensure your EBS R12.2 instance is supported

See [Source EBS R12.2 Instance Requirements](#) to ensure you can link your EBS R12.2 instance to the Delphix Engine.

Ensure your EBS 12.2 environments comply with Oracle's documentation

Your environments must comply with Oracle's requirements for installing EBS. These requirements are outlined on **Oracle E-Business Suite Release 12 Installation Guidelines (Doc ID 405565.1)** found at <https://support.oracle.com>.

Prepare the dbTier for linking

Delphix Engine's Unix Environment Requirements

The dbTier must meet the source requirements outlined in [Requirements for Unix Environments](#). These requirements are generic to all source Unix environments added to the Delphix Engine.

`oracle` User

The Delphix Engine must have access to an `oracle` user on the dbTier.

- This user should be a member of both the EBS `dba` and `oinstall` groups.
- The user should own the stage directory for DBTechStack and database files that will be cloned.
- Ensure that the "expect" utility exists on the remote host.

dbTechStack Binary Permissions

Verify that the `oracle` user described above has read permissions at the group level for:

- `$ORACLE_HOME/bin/nmb`
- `$ORACLE_HOME/bin/nmhs`
- `$ORACLE_HOME/bin/nmo`

The Delphix Engine's Oracle Database Requirements

The dbTier must meet the source requirements outlined in [Oracle Support and Requirements](#). These requirements are generic to all Unix environments containing an Oracle database to be linked.

Prepare the appsTier for Linking

Delphix Engine's Unix Environment Requirements

The appsTier must meet the source requirements outlined in [Requirements for Unix Environments](#). These requirements are generic to all source Unix environments added to the Delphix Engine.

`applmgr` User

The Delphix Engine must have access to an `applmgr` user on the appsTier.

- This user should be a member of the EBS `oinstall` group.
- The user should own the stage directory for Appstier files that will be cloned.

delphix_os OS user account

In order to separate authentication and perform privileged operations with a non-privileged OS account, first create an OS user account (i.e. "delphix_os") on the EBS DB Tier and appsTier node to be used as a source.

This user is easily created by the `createDelphixOSUser.sh` (located below on this page) script.

- The primary OS group of the Delphix Engine software owner account's (i.e. `delphix_os`) should be the same as the EBS Database/AppsTier software owner account (i.e. `oravis` or `applmgr`).
- Primary group = Oracle Install Group (typically `oinstall`), secondary group = OSDB Group (typically `dba`). There are lots of cases where the OS group named `dba` fills this role, so be sure to check the group membership of the EBS Database and AppsTier software owner account.
- Please note, the non-privileged OS account must have the **same** group as assigned to EBS Database or AppsTier privileged account (like `oravis` or `applmgr`).

Host Requirements:

To accomplish necessary tasks on the EBS Database and appsTier source hosts, the Delphix OS user account (henceforth referred to as "delphix_os") requires privilege elevation specifications.

Here is an example specification for the "sudo" privilege elevation utility, using the "visudo" to edit the "sudoers" configuration file. This specification makes the following assumptions:

- OS = Linux/Solaris
- OS account owning Oracle EBS Database Tier is named **oravis**
- OS account owning Oracle EBS appsTier is named **applvis**
- EBS Database and appsTier system base directory is **/u01/oracle/VIS**. This can be prefixed before all the below-mentioned commands to ensure the same are being executed from relevant paths.

 The following sudoers entry is only for template purpose. Modify the path in the below sudoers entry with the appropriate binary paths of your environment.

Entries required for Linking via low privileged user (delphix_os):

EBS DB Tier	<p>Defaults:delphix_os !requiretty</p> <pre>delphix_os ALL=NOPASSWD: /bin/su - oravis -c echo *, /bin/su - oravis -c rm -f */dlpx_force_autoflush*, /bin/su - oravis -c cp -f */ dlpx_force_autoflush*; chmod 755 */dlpx_force_autoflush*, /bin/su - oravis -c rm -rf */appsutil/clone/dbts*, /bin/su - oravis -c export PERL5LIB*perl -mdlpx_force_autoflush */adpreclone.pl dbTier*, /bin/su - oravis -c */rsync*, /bin/su - oravis -c test*mkdir*, /bin/su - oravis -c test*touch*chmod 750*cat*, /bin/su - oravis -c */EBS_kill/ kill_script.sh*, /bin/su - oravis -c rm -f */test_status.tmp*, /bin/ps, /bin/ su - oravis -c cp *dlpx_force_autoflush.pm*, /bin/su - oravis -c chmod 755 */dlpx_force_autoflush*</pre>
EBS appsTier	<p>Defaults:delphix_os !requiretty</p> <pre>delphix_os ALL=NOPASSWD: /bin/su - applvis -c echo *, /bin/su - applvis -c rm *.dlpx_run_edition*, /bin/su - applvis -c rm -f */ dlpx_force_autoflush*, /bin/su - applvis -c cd *echo *dlpx_force_autoflush*, /bin/su - applvis -c export PATH* export PERL5LIB* cd *perl -mdlpx_force_autoflush ./adpreclone.pl appsTier*, / bin/su - applvis -c */rsync*, /bin/su - applvis -c test*mkdir*, /bin/su - applvis -c test*touch*chmod 750*cat*, /bin/su - applvis -c */EBS_kill/ kill_script.sh*, /bin/su - applvis -c rm -f */test_status.tmp*, /bin/su - applvis -c *.env* sqlplus -s apps*, /bin/su applvis -c *.env* sqlplus -s apps*, /bin/su - applvis -c */*.env*echo*adadminsrvctl.sh status -nopromptmsg*, /bin/ps, /bin/su - applvis -c cp *dlpx_force_autoflush.pm*, /bin/su - applvis -c export PERL5LIB* cd *perl -mdlpx_force_autoflush */adpreclone.pl appsTier*</pre>

Requirement for Privilege Elevation Script: DLPX_DB_EXEC

In order to elevate privileges from a non-privileged OS account (like delphix_os) to a privileged OS account (like applmgr), we need to push a privilege elevation script (*dlpx_db_exec*) up into the Delphix virtualization engine to become part of the Delphix common plugin.

Why we need DLPX_DB_EXEC

Some customers want to use low privilege users to perform delphix operations like linking and provisioning. It means their low privilege user should have sudo permissions to execute EBS application and DB related commands. The privilege elevation profile script *dlpx_db_exec* allows them to execute commands that require superuser privileges on customer source and target machines.

The privilege elevation script *dlpx_db_exec* can be created or pushed to Delphix Engine using Web API calls, CURL or dxtoolkit.

For steps on creating a Privilege Elevation Profile please refer to [CLI Cookbook: How to Create or Edit a Privilege Elevation Profiles and profile Scripts](#)

Content of DLPX_DB_EXEC Privilege Elevation Profile:

```
#!/bin/sh
#
# Copyright (c) 2018 by Delphix. All rights reserved.
#
# This script allows customization of command execution with an alternate user
# account.
# Arg $1 contains "-u<optional user account>" for the desired user under
# which database commands will be executed.
# By default this argument is ignored and the script is executed as the
# default
# account.
#
if [[ $1 != -u* ]]; then
    echo "Incorrect command line parameters, -u<optional user account> is
    required as the first parameter"
    exit 1
fi
user_id=`echo $1 | sed -e "s/^-u//"`
shift 1
if [[ $user_id != "delphix_os" ]]; then
    command=$(printf "%s " "$@" )
    sudo su - $user_id -c "$command"
else
    $@
fi
```

Below is an example of how we can push privilege elevation script “dlpx_db_exec” on a customer Delphix Engine:

- Create a session to Delphix Engine as Delphix os User:

```
curl -i -c cookies.txt -X POST -H "Content-Type:application/json" http://
<Delphix-Engine>/resources/json/delphix/session -d '{
  "version":{
    "minor":11,
    "major":1,
    "micro": 5,
    "type":"APIVersion"
  },
  "type":"APISession"}
```

Note: The API Version needs to be identified as per the Delphix Engine installed at the customer site.

- Login to Delphix Engine as Delphix OS User:

```
curl -i -c cookies.txt -b cookies.txt -X POST -H "Content-Type:application/json" http://<Delphix-Engine>/resources/json/delphix/login -d '{
  "password":"delphix",
  "type":"LoginRequest",
  "target":"DOMAIN",
  "username":"delphix_admin"
}'
```

- Push DLPX_DB_EXEC contents to Delphix Engine:

```
curl -i -b cookies.txt -X POST -H "Content-Type:application/json" http://
<Delphix-Engine>/resources/json/delphix/host/privilegeElevation/profileScript/
HOST_PRIVILEGE_ELEVATION_PROFILE_SCRIPT-7 -d '{
  "type": "HostPrivilegeElevationProfileScript",
  "contents": "#\n# Copyright (c) 2018 by Delphix. All rights reserved.
\n#\n#\n#\n# This script allows customization of command execution with an
alternate user\n# account.\n\nif [[ $1 != -u* ]]; then\n  echo \"Incorrect
command line parameters, -u<optional user account> is required as the first
parameter\"\n\n exit 1\n\nfi\n\nuser_id=`echo $1 | sed -e \"s/^-u//\"`\n\nshift
1\n\nif [[ $user_id != \"delphix_os\" ]]; then\n\ncommand=$(printf \"%s \" \"$@\")
\n\nsudo su - $user_id -c \"$command\"\n\nelse\n\n$@\n\nfi\n"
}'
```

Script

[createDelphixOSUser.sh](#)

Related Topics

- [Oracle Support and Requirements](#)
- [Requirements for Unix Environments](#)
- [Provisioning a Virtual EBS R12.2 Instance](#)

Linking a Source EBS R12.2 Instance

This topic describes the process of linking an EBS R12.2 instance and creating the necessary dSources.

Prerequisite

Prepare your source EBS R12.2 instance for linking by following the outline in [Preparing a Source EBS R12.2 Instance for Linking](#).

Procedure

Link the Oracle Database

1. Link the Oracle database used by EBS, as outlined in [Linking an Oracle Data Source](#).
2. Link the Oracle Pluggable database for 19c MT, as outlined in [Linking an Oracle Pluggable Database](#).

Link the EBS dbTechStack

Prerequisite

Upload the EBS 12.2 dbTechStack plugin and refresh the environments during the preparation of the source environment. For more information, see [Upgrade EBS Objects to 6.0](#).

For successful linking of DBTECHSTACK from the standby site, follow any one of the procedures below:

- Place the physical standby database in READ ONLY/ READ ONLY WITH APPLY mode and run linking of source DBTECHSTACK. Even during subsequent snapsync operations, you have to make sure that the standby database container & its PDB is in READ ONLY / READ ONLY WITH APPLY mode.

 Make sure that the PDB is in READ ONLY mode, also the two default EBS PDB services `ebs_<PDBNAME>` and `<PDBNAME>_ebs_patch` should be in READY STATE. Linking will fail, if the PDB is not in READ ONLY mode.

- Perform the linking of dbTechStack(ORACLE_HOME) from the primary source site which is in READ_WRITE mode.

- 
- Before DBTechStack linking, to ensure that the hook operation succeeds during database provision, the ETCC patch level should be the latest. In the configure clone 1 hook `adcfgclone dbconfig` will perform the ETCC patch level check and will fail if it's not the latest.
 - While linking using the EBS plugin, the plugin runs `adpreclone.pl` dbTier in the background, which asks for apps database user password and connects with standby DB. If database is in MOUNT state then linking fails, as apps schema will not be able to login into standby DB.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **source dbTier environment** containing the source dbTechStack.

Linking from RAC dbTier

If you are linking from a RAC dbTier, select the environment for a single running node of the RAC cluster.

5. Click the **Environment Details** tab.

6. If the `oracle` environment user described in [Preparing a Source EBS R12.2 Instance for Linking](#) is not already added to the Delphix Engine, add the user.
7. Click the **Databases** tab.
8. Click the **Plus** icon within **E-Business Suite R12.2 dbTechStack** section under installations.
9. **Add Database** dialog box pops to add source config for dbTechStack.
10. Provide a user-defined name to the `Database Name` field and Oracle base install directory in the `Path` field.
11. Click on the **Add** button to successfully add the source config.
12. Re-check the **E-Business Suite R12.2 dbTechStack** section under installations for the `Add dSource` option mapped to the source config that is added in the above step.
13. Click **Manage > Datasets**.
14. Click the **Plus** icon next to Datasets and select **Add dSource**.
15. In the **Add dSource wizard**, select the Linked dSource as the type.
16. Enter the **EBS-specific parameters** for your dbTechStack.
17. These parameter values will be used when `adpreclone.pl` is run.
Ensure that the **DB Tier Context Name uses** the short hostname.
18. Select an Environment user.
19. **Privileged OS Account (Optional)** field should contain a high privileged user when the low privileged user is being used for linking.
20. Provide source oracle home directory name relative to oracle base installation directory in Oracle Home input field.
21. In the **DB Tier** Context Name field, provide the source context name; To get the DB Tier context name on source side:
 - For Oracle 12.1.0.2 version,

```
cd $ORACLE_HOME
source SID_hostname.env file
echo $CONTEXT_NAME
```

Below is an example,

Consider VISDB as the database SID and tavsrc-db as the short hostname.

```
[oravis@tavsrc-db VIS]$ source /u01/oracle/VIS/12.1.0.2/VISDB_tavsrc-
db.env
[oravis@tavsrc-db VIS]$ echo $CONTEXT_NAME
VISDB_tavsrc-db
```

Source Environment
tavsrc-db.dlpxdc.co

Data Type
AppData

Environment User

oravis

Privileged OS Account (Optional)

This privileged unix username will be used to link EBS dbTechStack. The privileged unix username should begin with a lower case letter or an underscore, followed by lower case letters, digits, underscores, or dashes. They can end with a dollar sign. Leave this field blank if you do not want to use privilege elevation.

Oracle Home *

12.1.0.2

Path where ORACLE_HOME will be relative to the dataset home being linked.

DB Tier Context Name *

VISDB_tavsrc-db

This will be used to prepare the source environment for linking and taking snapshots.

APPS Password credentials

This will be used to prepare source environment for linking and taking snapshots.

Username and Password

User name

Password *

....

Paths to exclude

This will be used to specify relative path(s) to exclude while linking DBTechstack.

data

oralnventory

1.1 TUNESDB_TX

Timeflow Status Configuration

Source Policies Masking Hooks Custom

Privileged OS Account (Optional)

Oracle Home
product/12.1.0.2/DbHome_1

DB Tier Context Name
TUN1_tristexadbadm01

APPS Password

Paths to exclude

checkpoints

log

tfa

audit

AUDIT_DIR

agent1310

patchdepot

cfgtoollogs

admin

diag

export

maintain

- For multi-tenant database, CONTEXT_NAME will be <PDB_NAME>_<short hostname>
- For RAC, CONTEXT_NAME will be <INSTANCE_SID>_<HOSTNAME>
For example: VIS2_exaukdb02

i When preparing the source Oracle RAC system for cloning, make sure that the linking is performed against the primary database node.

22. Provide apps schema password in the **Apps Password Credentials** section by selecting the **Username** and **Password** from the dropdown, leaving the username field empty.
23. Exclude the EBS database's data files if they are stored underneath the Oracle base install directory.
24. These data files will be linked with the database instead of with the dbTechStack. Add the "relative path" to the data files to the **Paths to Exclude** list. If the audit files (located under \$ORACLE_HOME/rdbms/audit) are large in size, the path of these files could also be excluded so as to avoid dbTechStack timeout issues during provisioning.

i Ensure that you add the relative paths and NOT the absolute paths.

For example, if you want to exclude the `/u01/oracle/VIS/data` path, then add `"data"` only under the **Paths to Exclude** list. Here, `/u01/oracle/VIS/data` is the absolute path and `"data"` is the relative path to be added.

25. Click **Next**.
26. Enter a **dSource Name**.
27. Select a **Database Group** for the dSource.
28. Click **Next**.
Adding a dSource to a database group enables you to set Delphix Domain user permissions for that dSource's objects, such as snapshots.
29. Select a **SnapSync** policy.
30. Click **Next**.
31. Enter any **custom pre or post sync logic** as Pre-Sync or Post-Sync hook operations.
Remember that `adpreclone.pl dbTier` is already run prior to every Snapshot of the dbTechStack.
The Pre-Sync hook operations will be run prior to running the `adpreclone.pl` tool.
32. Click **Next**.
33. Review the **dSource Configuration** and **Data Management** information, and then click **Finish**.
The Delphix Engine will initiate two jobs to create the dSource, **DB_Link** and **DB_Sync**. You can monitor these jobs by clicking **Active Jobs** in the top menu bar, or by selecting **System > Event Viewer**. When the jobs have completed successfully, the **files** icon will change to a **dSource** icon on the **Environments > Databases** screen, and the dSource will be added to the list of **Datasets** under its assigned group.

Link the EBS appsTier

Prerequisite

Upload the EBS 12.2 dbTechStack plugin and refresh the environments during the preparation of source environment. For more information, see [Upgrade EBS Objects to 6.0](#).

1. Log in to the **Delphix Management** application using **Delphix Admin** credentials.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **source appsTier environment**.



Linking from multi-node appsTier

If you are linking from a multi-node appsTier, select the environment for the node on which EBS admin services reside.

5. Click the **Environment Details** tab.
6. If the `app\mgr` environment user described in [Preparing a Source EBS R12.2 Instance for Linking](#) is not already added to the Delphix Engine, add the user.
7. Click the **Databases** tab.
8. Click the **Plus** icon within **E-Business Suite R12.2 appsTier** section under installations.
9. **Add Database** dialog box pops to add source config for dbTechStack.
10. Provide a user-defined name to the ``Database Name`` field and Oracle base install directory in the ``Path`` field.
11. Click on the **Add** button to successfully add the source config.
12. Re-check the **E-Business Suite R12.2 appsTier** section under installations for the **Add dSource** option mapped to the source config that is added in the above step.
13. Click **Manage > Datasets**.
14. Click the **Plus** icon next to Datasets and select **Add dSource**.

15. In the **Add dSource wizard**, select the **Linked dSource** as the type.
16. In the **Add dSource wizard**, select the **appsTier files source** you just created.
17. Enter the **EBS-specific parameters** for your appsTier.

These parameter values will be used when `adpreclone.pl` is run.



18. Select an Environment User.
19. **Privileged OS Account (Optional)** field should contain a high privileged user when the low privileged user is being used for linking.
20. Provide the apps schema password in the **Apps Password** section by choosing the **Username** and **Password** from the dropdown, leaving the username field empty.
21. Provide the weblogic password in the **Weblogic AdminServer Password** section by choosing the **Username** and **Password** from the dropdown, leaving the username field empty.
22. Add the "relative paths" of files to exclude to the **Paths to Exclude** list.

23. Exclude the non-edition file system, instance directories, and FMW home directories which will be recreated when provisioning a virtual EBS appsTier.

These paths should typically include:

- `fs_ne`
- `fs1/inst`
- `fs2/inst`
- `fs1/FMW_Home`
- `fs2/FMW_Home`
- `fs2/EBSapps` or `fs1/EBSapps` depending on patch FS in your environment.
- `EBSapps.env`

You can exclude all other directories or files placed in `$APPS_BASE` path.

 Ensure that you add the relative paths and NOT the absolute paths. For example, if you want to exclude the `/u01/oracle/VIS/fs1/inst` path, then add `"fs1/inst"` only under the **Paths to Exclude** list. Here, `/u01/oracle/VIS/fas1/inst` is the absolute path and `"fs1/inst"` is the relative path to be added.

24. Click **Next**.
25. Enter a **dSource Name**.
26. Select a **Database Group** for the dSource.
27. Click **Next**.
Adding a dSource to a database group enables you to set Delphix Domain user permissions for that dSource's objects, such as snapshots.
28. Click **Next**.
29. Select a **SnapSync** policy.
30. Click **Next**.
31. Enter any **custom pre- or post-sync logic** as Pre-Sync or Post-Sync hook operations.
Remember that `adpreclone.pl` appsTier is already run prior to every SnapSync of the appsTier.
The Pre-Sync hook operations will be run prior to running the `adpreclone.pl` tool.
32. Click **Next**.
33. Review the **dSource Configuration** and **Data Management** information, and then click **Finish**.
The Delphix Engine will initiate two jobs to create the dSource, **DB_Link**, and **DB_Sync**. You can monitor these jobs by clicking **Active Jobs** in the top menu bar, or by selecting **System > Event Viewer**. When the jobs have completed successfully, the **files** icon will change to a **dSource** icon on the **Environments > Databases** screen, and the dSource will be added to the list of **Datasets** under its assigned group.

Related Topics

- [Oracle Support and Requirements](#)
- [Requirements for Unix Environments](#)
- [Preparing a Source EBS R12.2 Instance for Linking](#)

Preparing Target EBS R12.2 Environments for Provisioning

This topic outlines the prerequisites for provisioning a virtual EBS R12.2 instance to target environments.

Ensure your target EBS R12.2 instance is supported

See [Virtual EBS R12.2 Instance Requirements](#) to ensure you can provision your EBS R12.2 instance.

Ensure your EBS 12.2 environments comply with Oracle's documentation

Your environments must comply with Oracle's requirements for installing EBS. These requirements are outlined on **Oracle E-Business Suite Release Notes, Release 12.2 (Doc ID 1320300.1)** found at <https://support.oracle.com>.

 Oracle has released an E-Business Suite Pre-Install RPM (available on ULN and public yum) that includes all required RPMs for both the appsTier and dbTier of an R12.2 installation. Details can be found in **Oracle E-Business Suite Installation and Upgrade Notes Release 12 (12.2) for Linux x86-64 (Doc ID 1330701.1)** found at <https://support.oracle.com>.

Prepare the dbTier for provisioning

Delphix Engine's Unix Environment Requirements

The dbTier must meet the target requirements outlined in [Requirements for Unix Environments](#). These requirements are generic to all target Unix environments you add to the Delphix Engine.

`oracle` User

The Delphix Engine must have access to an `oracle` user on the dbTier.

- This user should be a member of both the EBS `dba` and `oinstall` groups.
- This user will be given proper permissions to manage the dbTechStack and database.
- Ensure that the "expect" utility exists on the remote host.

`oraInst.loc`

An `oraInst.loc` file must exist on the dbTier prior to provisioning. This file will specify where the `oraInventory` directories live or where they should be created if they do not already exist.

The `oraInst.loc` file is typically located at `/etc/oraInst.loc` on Linux or `/var/opt/oracle/oraInst.loc` on Solaris. Ensure that the `oraInventory` to which this file points is writeable by the `oracle` user.

Consult Oracle EBS documentation for more information about where to place this file on your dbTier and what this file should contain.

Oracle Central Inventory

The Oracle Central Inventory is a repository for all Oracle products (software) installed on a system. Since the Central Inventory consists of system-specific information, it is required that the Central Inventory be saved on a local non-shared directory on the system. While software can be shared across nodes, the inventory should be local to each system. The global inventory/central inventory(`oraInventory`) for each node should be on a local file system other than Delphix mount base path.

For more info, see the following notes

- **Note 360079.1- Global and Local Inventory explained.**
- **Note 564192.1- FAQs on Central Inventory and Oracle Home Inventory (Local Inventory) in Oracle RDBMS.**
- **Note 834894.1- How to Create a Clean oraInventory in Release 12?**
- **Note 742477.1 - How to create, update or rebuild the Central Inventory for Applications R12?**
- **Note 295185.1- How to Recreate the Global oraInventory?**

/etc/hosts

Use the following to verify hostname settings:

For Oracle Linux 5, 6, and 7, and Red Hat Enterprise Linux 5, 6, and 7:

Verify that the /etc/hosts file is formatted as follows:

```
127.0.0.1 localhost.localdomain localhost
[ip_address] [node_name].[domain_name] [node_name]
```

 Hostname and its IP address of Database host and Appstier host(s) (primary and secondary hosts for multinode) should be added in the `/etc/hosts` file of all the Database, host, Appstier host(s) (primary and secondary hosts for Multinode).

These correct entries are required to avoid adcfgclone timeout issue on the prompt for "**Target System Domain Name**".

Clean Up Before Provisioning Option

If you plan to utilize the Clean up Before Provision option available during dbTechStack provisioning, the Delphix Engine requires the Database Oracle Home to be patched with Oracle Universal Installer (OUI) version 10.2 or above. You can read more about the Clean Up Provisioning option in Provisioning a Virtual EBS R12.2 Instance. Note that provisioning is still possible without this option specified, but you will need to manage the target dbTier's Oracle Inventory manually to ensure that conflicting entries do not cause provisions to fail.

The Delphix Engine's Oracle Database Requirements

The dbTier must meet the target requirements outlined in [Oracle Support and Requirements](#). These requirements are generic to all target Unix environments expected to host a virtual Oracle database.

Prepare the appsTier for Provisioning

Delphix Engine's Unix Environment Requirements

The appsTier must meet the target requirements outlined in [Requirements for Unix Environments](#). These requirements are generic to all target Unix environments added to the Delphix Engine.

`appImgr` User

The Delphix Engine must have access to an `appImgr` user on the appsTier.

- This user should be a member of the EBS `oinstall` group.
- This user will be given proper permissions to manage the appsTier.

delphix_os OS user account

In order to separate authentication and perform privileged operations with a non-privileged OS account, first create an OS user account (i.e. "delphix_os") on the EBS DB Tier and appsTier node to be used as a target. This user is easily created by the createDelphixOSUser.sh (located below on this page) script.

- The primary OS group of the Delphix Engine software owner account's (i.e. delphix_os) should be the same as the EBS Database/AppsTier software owner account (i.e. oravis or applmgr).
- Primary group = Oracle Install Group (typically oinstall), secondary group = OSDBA Group (typically dba). There are lots of cases where the OS group named dba fills this role, so be sure to check the group membership of the EBS Database and AppsTier software owner account.
- Please note, the non-privileged OS account must have the **same** group as assigned to EBS Database or appsTier privileged account (like oravis or applmgr).
- Ensure that "expect", "wget", and "lsf" utility exist on the remote host.

Host Requirements for Multinode:

- The primary host and the secondary host user name and its UID, Group name and GID should be same to avoid the permission issue during the provisioning operation.
- Check the high and low privilege user id and GID in the primary and secondary node and change it accordingly using groupmod with the following commands:

```
$ id oravis
uid=509(oravis) gid=502(oinstall)
$ id delphix_os
uid=510(delphix_os) gid=502(oinstall)
Change the group id as root using,
...
...
# groupmod -g 502 oinstall
# usermod -g oinstall oravis
# usermod -g oinstall delphix_os
```

 Primary and secondary appstier host need to be configured with passwordless ssh if the target WLS password is going to be different from the source WLS password. Passwordless ssh configuration on the Appstier multinode host requirement is required for configuring WLS password different from source.

Host Requirements:

To accomplish necessary tasks on the EBS Database and appsTier target hosts, the Delphix OS user account (henceforth referred to as "delphix_os") requires privilege elevation specifications. Ensure that "wget" and "lsf" utilities exist on the remote host.

Here is an example specification for the "sudo" privilege elevation utility, using the "visudo" to edit the "sudoers" configuration file. This specification makes the following assumptions:

- OS = Linux
- OS account owning Oracle EBS Database Tier is named **oravis**
- OS account owning Oracle EBS appsTier is named **applvis**
- EBS Database and appsTier system base directory is **/u01/oracle/VIS**. This can be prefixed before all the below-mentioned commands to ensure the same is being executed from relevant paths.

 The following sudoers entry is only for template purposes. Modify the path in the below sudoers entry with the appropriate binary paths of your environment.

Entries required for Provisioning via low privileged user (delphix_os):

EBS DB Tier	<p>Defaults:delphix_os !requiretty</p> <pre>delphix_os ALL=NOPASSWD: /bin/su - oravis -c /bin/mount, /bin/su - oravis -c /bin/umount, /bin/su - oravis -c echo *, /bin/su - oravis -c export *, /bin/ su - oravis -c test*mkdir*, /bin/su - oravis -c test*touch*chmod 750*cat*, / bin/su - oravis -c */EBS_kill/kill_script.sh*, /bin/su - oravis -c rm -f */ test_status.tmp*, /bin/su - oravis -c cp -f */dlpx_force_autoflush*; chmod 755 */dlpx_force_autoflush*, /bin/su - oravis -c rm -f */ dlpx_force_autoflush*, /bin/su - oravis -c */*.env*bin/lsnrctl status*, /bin/ su - oravis -c */*.env*/addlnctl.sh*, /bin/su - oravis -c rm -f /u01/oracle/VIS/ */.delphix_adpreclone.lck*, /bin/su - oravis -c rm -rf */appsutil/clone/ dbts*, /bin/su - oravis -c touch /u01/oracle/VIS/ */.delphix_adpreclone.lck*, /bin/su - oravis -c */bin/runInstaller -silent -detachHome*, /bin/su - oravis -c export PATH* export PERL5LIB*perl -mdlpx_force_autoflush */adcfgclone.pl dbTechStack*, /bin/su - oravis -c export PERL5LIB*perl -mdlpx_force_autoflush */adpreclone.pl dbTechStack*, /bin/su - oravis -c cd *; make -f ins_rdbms.mk*ioracle*, /bin/ su - oravis -c cd *; make -f ins_rdbms.mk dnfs_off*, /bin/su - oravis -c rm -f *bak, /bin/su - oravis -c mv *, /bin/su - oravis -c */*.env*sqlplus* as sysdba*, /bin/su - oravis -c */*.env*perl */appsutil/clone/bin/adcfgclone.pl dbconfig*, /bin/su - oravis -c */*.env*perl */appsutil/scripts*/adpreclone.pl database*, /bin/su - oravis -c sed*sqlnet.ora*, /bin/su - oravis -c */ *.env*sqlplus apps*, /bin/su - oravis -c */*.env*; make -f *rdbms/lib/ ins_rdbms.mk dnfs_off*, /bin/su - oravis -c */*.env*; ln -s *, /bin/mount, / bin/umount, /bin/ps, /bin/mkdir, /bin/su - oravis -c source* -dboraclehome* perl*txkGenCDBTnsAdmin.pl*, /bin/su - oravis -c */ *.env*perl */appsutil/bin/txkCfgUtlfileDir.pl -contextfile*, /bin/su - oravis -c */*.env* mkdir -p*, /bin/su - oravis -c chmod 775 *dbs*, /bin/su - oravis -c chmod 6751 */bin/oracle*, /bin/su - oravis -c cp *dlpx_force_autoflush.pm*, /bin/su - oravis -c chmod 755 */ dlpx_force_autoflush*, /bin/su - oravis -c umask*touch *source_apps_file.txt, /bin/su - oravis -c mkdir -p*, /bin/su - oravis -c cp *pairsfile*, /bin/su - oravis -c *perl -mdlpx_force_autoflush */ adclonectx.pl*, /bin/su - oravis -c *perl -mdlpx_force_autoflush */ adcfgclone.pl dbTechStack*, /bin/su - oravis -c touch */.delphix_adpreclone.lck*, /bin/su - oravis -c rm -f */.delphix_adpreclone.lck*</pre>
EBS appsTier	Defaults:delphix_os !requiretty

EBS DB Tier	<p>Defaults:delphix_os !requiretty</p> <pre>delphix_os ALL=NOPASSWD: /bin/su - oravis -c /bin/mount, /bin/su - oravis -c /bin/umount, /bin/su - oravis -c echo *, /bin/su - oravis -c export *, /bin/ su - oravis -c test*mkdir*, /bin/su - oravis -c test*touch*chmod 750*cat*, / bin/su - oravis -c */EBS_kill/kill_script.sh*, /bin/su - oravis -c rm -f */ test_status.tmp*, /bin/su - oravis -c cp -f */dlpx_force_autoflush*; chmod 755 */dlpx_force_autoflush*, /bin/su - oravis -c rm -f */ dlpx_force_autoflush*, /bin/su - oravis -c */*.env*bin/lsnrctl status*, /bin/ su - oravis -c */*.env*/addlnctl.sh*, /bin/su - oravis -c rm -f /u01/oracle/VIS/ */.delphix_adpreclone.lck*, /bin/su - oravis -c rm -rf */appsutil/clone/ dbts*, /bin/su - oravis -c touch /u01/oracle/VIS/ */.delphix_adpreclone.lck*, /bin/su - oravis -c */bin/runInstaller -silent -detachHome*, /bin/su - oravis -c export PATH* export PERL5LIB*perl -mdlpx_force_autoflush */adcfgclone.pl dbTechStack*, /bin/su - oravis -c export PERL5LIB*perl -mdlpx_force_autoflush */adpreclone.pl dbTechStack*, /bin/su - oravis -c cd *; make -f ins_rdbms.mk*ioracle*, /bin/ su - oravis -c cd *; make -f ins_rdbms.mk dnfs_off*, /bin/su - oravis -c rm -f *bak, /bin/su - oravis -c mv *, /bin/su - oravis -c */*.env*sqlplus* as sysdba*, /bin/su - oravis -c */*.env*perl */appsutil/clone/bin/adcfgclone.pl dbconfig*, /bin/su - oravis -c */*.env*perl */appsutil/scripts*/adpreclone.pl database*, /bin/su - oravis -c sed*sqlnet.ora*, /bin/su - oravis -c */ *.env*sqlplus apps*, /bin/su - oravis -c */*.env*; make -f *rdbms/lib/ ins_rdbms.mk dnfs_off*, /bin/su - oravis -c */*.env*; ln -s *, /bin/mount, / bin/umount, /bin/ps, /bin/mkdir, /bin/su - oravis -c source* -dboraclehome* perl*txkGenCDBTnsAdmin.pl*, /bin/su - oravis -c */ *.env*perl */appsutil/bin/txkCfgUtlfileDir.pl -contextfile*, /bin/su - oravis -c */*.env* mkdir -p*, /bin/su - oravis -c chmod 775 *dbs*, /bin/su - oravis -c chmod 6751 */bin/oracle*, /bin/su - oravis -c cp *dlpx_force_autoflush.pm*, /bin/su - oravis -c chmod 755 */ dlpx_force_autoflush*, /bin/su - oravis -c umask*touch *source_apps_file.txt, /bin/su - oravis -c mkdir -p*, /bin/su - oravis -c cp *pairsfile*, /bin/su - oravis -c *perl -mdlpx_force_autoflush */ adclonectx.pl*, /bin/su - oravis -c *perl -mdlpx_force_autoflush */ adcfgclone.pl dbTechStack*, /bin/su - oravis -c touch */.delphix_adpreclone.lck*, /bin/su - oravis -c rm -f */.delphix_adpreclone.lck*</pre>
	<pre>delphix_os ALL=NOPASSWD: /bin/su - applvis -c echo *, /bin/su - applvis -c rm *.dlpx_run_edition*, /bin/su - applvis -c rm -f *dlpx_force_autoflush*, / bin/su - applvis -c cd *echo *dlpx_force_autoflush*, /bin/su - applvis -c export PATH* export PERL5LIB* cd *perl -mdlpx_force_autoflush ./ adpreclone.pl appsTier*, /bin/su - applvis -c */rsync*, /bin/su - applvis -c test*mkdir*, /bin/su - applvis -c test*touch*chmod 750*cat*, /bin/su - applvis -c */EBS_kill/kill_script.sh*, /bin/su - applvis -c rm -f */ test_status.tmp*, /bin/su - applvis -c */*.env* sqlplus -s *apps*, /bin/su - applvis -c */adstrtal.sh*, /bin/su - applvis -c */adstpall.sh*, /bin/su - applvis -c */*.env* cd *adautoCfg.sh*, /bin/su - applvis -c export PATH* export PERL5LIB* cd *perl -mdlpx_force_autoflush ./adcfgclone.pl appsTier*, /bin/ su - applvis -c *perl -mdlpx_force_autoflush */adclonectx.pl addnode contextfile*pairsfile*outfile*, /bin/su - applvis -c *adadminsrvctl.sh*, /bin/ su - applvis -c *admanagersrvctl.sh*, /bin/su - applvis -c *adnodemgrctl.sh*, /bin/su - applvis -c */*.env* cd *perl -mdlpx_force_autoflush ./txkUpdateEBSDomain.pl*contextfile*action*upd ateAdminPassword*, /bin/su - applvis -c */bin/runInstaller -silent -detachHome*, /bin/su - applvis -c rm -rf */inst/apps/*, /bin/su - applvis -c rm -rf *FMW_Home*, /bin/su - applvis -c rm -rf *fs*, /bin/su - applvis -c find*exec rm -rf *, /bin/su - applvis -c cp */inst/apps/*appl/admin */inst/ apps/*, /bin/su - applvis -c */*EBSapps.env*perl */patch/115/bin/ txkSetAppsConf.pl *contextfile*configoptio*n*oacore*oafm*forms*formsc4ws*, /bin/su - applvis -c rsync -aH --delete --ignore-errors */EBSapps/ */EBSapps/*, /bin/</pre>

oraInst.loc

An oraInst.loc file must exist on every appsTier node prior to provisioning. This file will specify where the oraInventory directories live or where they should be created if they do not already exist.

The oraInst.loc file is typically located at `/etc/oraInst.loc` on Linux or `/var/opt/oracle/oraInst.loc` on Solaris. Ensure that the oraInventory to which this file points is writable by the `appMgr` user or application tier OS user.

If you are provisioning a single-node appsTier,

- [Required] The global inventory/central inventory(oraInventory) for each node should be on a local file system other than Delphix mount base path.

If you are provisioning a multi-node appsTier,

- The following recommendations are mandatory;
 - While configuring SHARED APPL_TOP application, Delphix expects the oraInventory directory path should be beneath the Delphix mount base path.
 - The `inventory_loc` variable defined in `file /etc/oraInst.loc` on all application servers should point to same shared oraInventory directory location on shared Delphix mount base path.
 - The Delphix Engine's cloning workflow requires that all nodes in the appsTier have access to the shared oraInventory directories via Delphix-provided storage.
- [Required] The Delphix Engine's automation requires that all nodes in the appsTier have access to the oraInventory directories via Delphix-provided storage.

Objective of placing oraInventory on Shared Delphix mount storage

While configuring shared APPL_TOP application, instead of storing oraInventory on the local directory of each server, all application servers should point to the same oraInventory directory on shared Delphix mount storage, so that you need to patch only on a single server and the oraInventory gets in-sync for all application servers thus preventing the oraInventory of remaining servers becoming invalid.

 Details can be found in **Sharing The Application Tier File System in Oracle E-Business Suite Release 12 (Doc ID 384248.1)** found at <https://support.oracle.com>. You can also consult Oracle EBS documentation for more information about where to place this file on your appsTier and what this file should contain.

Clean Up Before Provisioning Option

If you plan to utilize the Cleanup Before Provision option available during appsTier provisioning, the Delphix Engine requires the Tools Oracle Home to be patched with Oracle Universal Installer (OUI) version 10.2 or above. You can read more about this provisioning option in Provisioning a Virtual EBS R12.2 Instance.

Note that provisioning is still possible without this option specified, but you will need to manage the target appTier's Oracle Inventory manually to ensure that conflicting entries do not cause provisions to fail.

Script

[createDelphixOSUser.sh](#)

Related Topics

- [Requirements for Unix Environments](#)
- [Oracle Support and Requirements](#)
- [Virtual EBS R12.2 Instance Requirements](#)

Provisioning a Virtual EBS R12.2 Instance

This topic describes the process of provisioning a virtual EBS R12.2 instance.

- ⚠️ Maintain the integrity in names(CDB name, PDB name, DB hostname, /etc/hosts file format nomenclature). During dbTechStack, vPDB or vDB provisioning, maintain the same lower case and use the same name vPDB during virtual Appstier provision as well. There is no restriction on using UPPER or LOWER case characters. However, it is recommended to maintain consistency.

Prerequisites

- You must have already linked a source instance of EBS R12.2. For more information, see [Linking a Source EBS R12.2 Instance](#).
- Prepare your target EBS R12.2 environments for provisioning by following the outline in [Preparing Target EBS R12.2 Environments for Provisioning](#).

Snapshot Coordination

Changes applied to EBS and picked up only in certain dSource snapshots may make certain combinations of snapshots across the appsTier and dbTier incompatible. When provisioning, refreshing or rewinding a virtual EBS instance, be sure the points in time you select for each dataset are compatible with each other.

- ⚠️ The source DB host may have multiple DATA_TOPs but the target DB host will be provisioned with a single DATA_TOP that resides on the target Mount Base path provided during the database provisioning.

Provisioning the EBS dbTechStack

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **dbTechStack dSource**.
5. Click the **TimeFlow** tab.
6. Select a dSource **snapshot**.
7. Click **Provision**.
The **Provision VDB** wizard will open.
8. Select an **Environment**.
This environment will host the virtual dbTechStack and be used to execute hook operations specified in step 15 in the **s** section.
9. Select **E-Business Suite R12.2 dbTechStack** from the **Installation** dropdown.
10. Select an **Environment User**.
11. This user should be the `oracle` user-outlined in [Preparing Target EBS R12.2 Environments for Provisioning](#).
12. Enter a **Mount Path** for the virtual dbTechStack files.
13. Enter the **EBS-specific parameters** for the virtual dbTechStack. A subset of these parameters is discussed in more detail below.
 - a. The **Privileged OS Account (Optional)** field should contain the high privileged user when the low privileged user is being used for provisioning.
 - b. The **Target APPS Password** is the new apps password that is required to configure virtual dbTechStack. This password is encrypted when stored within the Delphix Engine and is available as an environment variable to the `adcfgclone` process.

- c. Provide source oracle home directory name relative to oracle base installation directory that was provided during linking of dbTechStack in Oracle Home input field.
 - d. Ensure that the **Target DB Hostname** value is the short hostname, not the fully-qualified hostname.
 - e. The **Target DB/CDB SID** is the new database SID (CDB SID in case of 19c) that is required to configure virtual dbTechStack.
 - f. The **Target PDB Name** field is added for the pluggable database to be configured for Oracle 19c database. Please leave this field empty when using Oracle 11g/12c database.
 - g. Provide **Target utl_file_dir**, the default value for this field will be `"/var/tmp"`. The **utl_file_dir** initialization parameter is deprecated in Oracle Database 12c Release 2 (12.2.0.1), and maybe unsupported in a future release.
 - h. Provide **DISPLAY Variable**, the default value is `hostname:0.0`
 - i. Provide **Target Port Pool**, please provide a port pool that is available. By default, the value is 1. The range for this value is 0 to 99.
 - j. Enable the **Disable RAC** option if you want to permit the Delphix Engine to automatically disable the RAC option for the binaries when applicable.
This option is necessary if provisioning from a dSource with RAC dbTier because the binaries are relinked with the `rac_on` option even after running `adcfgclone`. If the source binaries already have the RAC option disabled (also the case for SI dbTier), the Delphix Engine ignores this option.
 - k. Enable the **Cleanup Before Provision** option if you want to permit the Delphix Engine to automatically clean up stale EBS configuration during a refresh. This option is recommended, but only available if your Oracle Home is patched with Oracle Universal Installer (OUI) version 10.2 or above.
 - i. With this option enabled, the Delphix Engine will inspect the target environment's oraInventory prior to refreshing this virtual dbTechStack. If any Oracle Homes are already registered within the specified **Mount Path**, the Delphix Engine will detach them from the inventory prior to running `adcfgclone`. These homes must be detached prior to running post-clone configuration. If they are not detached, `adcfgclone` will fail, citing conflicting oraInventory entries as an issue.
 - ii. Without this option enabled, Oracle Homes that conflict with the specified **Mount Path** will be reported in an error instead of automatically detached. For refresh to succeed, you must manually detach conflicting Oracle Homes prior to refresh.
 - l. Provide **Oracle OS User** and **Group**, it's optional.
14. Click **Next**.
 15. Enter a **VDB Name**.
 16. Select a **Target Group** for the VDB.
If necessary, click the green **Plus** icon to add a new group.
 17. Select a **Snapshot Policy** for the VDB.
If necessary, click the **Plus** icon to create a new policy.

Snapshot conflicts

When Snapshot is running against the dbTechStack, database, or appsTier, the Delphix Engine also executes pre-clone logic to ensure the latest configuration is staged in the captured snapshots. Unfortunately, if multiple Snapshots are running against the same EBS instance concurrently, this pre-clone logic may fail and produce bad snapshots.
To avoid SnapSync conflicts, spread out your SnapSync policies for an EBS instance by one hour or more.

18. Click **Next**.
19. Enter any **custom hook operations** that are needed to help correctly manage the virtual dbTechStack files. For more information about these hooks, when they are run, and how operations are written.
The Configure Clone hook will be run after the

```
adcfgclone.pl
```

tool has both mounted and configured the dbTechStack.

20. Click **Next**.
21. Click **Submit**.

When provisioning starts, you can review the progress of the job in the **Datasets** panel, or in the **Job History** panel of the **Dashboard**. When provisioning is complete, the dbTechStack VDB will be included in the group you designated and listed in the **Datasets** panel. If you select the dbTechStack VDB in the **Datasets** panel and click the **Configuration** tab, you can view information about the virtual files and its Data Management settings.

For tips on monitoring the progress of dbTechStack provisioning, see [Monitoring EBS R12.2 dbTechStack Provisioning Progress](#).

Provisioning the Oracle 19C Database

1. Refresh the Target DB environment from **Manage > Environments** so that the Listener brought up during dbTechStack provision gets discovered and the Installation Home to be listed in dropdown during Database provision. If Target DB environment is not refreshed, you will see `This environment has no compatible Oracle Installation Homes` in the Installation Home dropdown in the **Database provision**.
2. Provision the EBS database to the target dbTier environment by following the steps outlined in [Provisioning an Oracle VDB](#). For Oracle EBS database for 19c MT, provision the database by following the steps outlined in [Provisioning an Oracle Virtual Pluggable Database](#).

Snapshot conflicts

When Snapshot is running against the dbTechStack, database, or appsTier, the Delphix Engine also executes pre-clone logic to ensure the latest configuration is staged in the captured snapshots. Unfortunately, if multiple Snapshots are running against the same EBS instance concurrently, this pre-clone logic may fail and produce bad snapshots. To avoid SnapSync conflicts, spread out your SnapSync policies for an EBS instance by one hour or more.

3. Select the correct **Installation Home**.
This should be the virtual dbTechStack you just added to the Delphix Engine.
4. Select an Environment User.
5. This user should be the oracle user-outlined in Preparing Target EBS R12.2 Environments for Provisioning.
6. For the **Container Database** option, select the checkbox next to **Create a New Container Database**, and click **Next**.
7. Select a Target Group for the VDB.
8. In the vPDB **Configuration** section, provide the Target PDB name given during dbTechStack provision in the **Oracle Pluggable Database Name** field and provide any unique user-defined name in **vPDB Name** to be displayed in the Delphix Engine.
9. In vCDB **CONFIGURATION** section, provide the Target DB/CDB SID name given during dbTechStack provision in the **Database Name** field and SID and provide any unique user-defined name in the **vCDB Name** and **Database Unique Name** field, and click **Next**.
10. Click **Advanced**.
11. Select the correct **Oracle Node Listeners** value.
This should be the listener corresponding to the virtual dbTechStack you just added to the Delphix Engine.
12. Add the EBS R12.2 dbTier environment file as a **Custom Environment Variables** entry.
13. This file can be specified as an Environment File with Path Parameters of `$ORACLE_HOME/
<CONTEXT_NAME>.env` .
Replace `<CONTEXT_NAME>` with the virtual EBS instance's context name. The Delphix Engine will expand

the `$ORACLE_HOME` variable at runtime.

For more information, see [Customizing Oracle VDB Environment Variables](#).

14. Add a **Run Bash Shell Command** operation to the Configure Clone hook to ensure that `adcfgclone` is run against the newly provisioned database for a **high privileged user**, as shown in the script below.

 You can pass the credentials securely to **Hook Operations** by setting up the base variables. See [Passing Credentials Securely to Hook Operations](#) for more details.

Configure Clone First Hook for a 19c Multi-Tenant High Privileged User

```
#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.
CONTEXT_NAME=${DELPHIX_PDB_NAME}_${hostname -s}
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
. "${ORACLE_HOME}/${ORACLE_SID}_${hostname -s}.env";
# Check for local_listener parameter is set for PDB, otherwise set it
# appropriately
check_value=$(sqlplus -s "/" as sysdba" <<EOF
alter session set container="${DELPHIX_PDB_NAME}";
show parameter local_listener;
EOF
)
local_listener=$(echo "$check_value" | awk '{print $11}')
value=("${local_listener//:/ }")
host="${value[0]}"
port="${value[1]}"
curr_port=$(grep PORT < "${ORACLE_HOME}/network/admin/listener.ora" | awk
'{print $9}' | sed 's/)//g')
if [[ $port != "$curr_port" || $host != "$(hostname -s)" ]]; then
sqlplus -s "/" as sysdba" <<EOF
alter session set container="${DELPHIX_PDB_NAME}";
alter system set local_listener='${hostname -s}:${curr_port}';
alter system register;
EOF
fi
sqlplus "/" as sysdba" <<EOF
@${ORACLE_HOME}/appsutil/install/${CONTEXT_NAME}/adupplib.sql so
EOF
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
perl "${ORACLE_HOME}/appsutil/clone/bin/adcfgclone.pl" dbconfig "${
ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml" <<EOF # noqa
${APPS_PASSWD}
EOF
```

15. Configure Clone hook to ensure that `adcfgclone` is run against the newly provisioned database for a **low privileged user**, as shown in the script below.

 Customers have to identify the `DLPX_DB_EXEC_SCRIPT` value from their remote environment, this would always be inside the toolkit directory provided while adding the environment on Delphix Engine.

Configure Clone First Hook for a 19c Multi-Tenant Low Privileged User

```
#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.
DLPX_DB_EXEC_SCRIPT="<Remote BIN location till dlp_x_db_exec script>"
DLPX_PRIV_USER=oravis
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
CONTEXT_NAME=${DELPHIX_PDB_NAME}_$(hostname -s)

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${ORACLE_SID}
_${hostname -s}.env;"

# Check for local_listener parameter is set for PDB, otherwise set it
appropriately
check_value=$(("${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/
${CONTEXT_NAME}.env; sqlplus -s \"/ as sysdba\" <<-EOF
alter session set container=${DELPHIX_PDB_NAME};
show parameter local_listener;
EOF
")

local_listener=$(echo "$check_value" | awk '{print $11}')
value=("${local_listener//:/ }")
host="${value[0]}"
port="${value[1]}"

curr_port=$(grep PORT < "${ORACLE_HOME}/network/admin/listener.ora" | awk
'{print $9}' | sed 's/)//g')

if [[ $port != "$curr_port" || $host != "$(hostname -s)" ]]; then
  "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{ORACLE_SID}_$(hostname -s).env; sqlplus -s \"/ as sysdba\" <<-EOF
  alter session set container=${DELPHIX_PDB_NAME};
  alter system set local_listener='${hostname -s}:${curr_port}';
  alter system register;
EOF
"
```

```

fi
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${ORACLE_SID}
_$(hostname -s).env; sqlplus \"/ as sysdba\" <<-EOF
@${ORACLE_HOME}/appsutil/install/${CONTEXT_NAME}/adupplib.sql so
EOF
"
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${
CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/clone/bin/adcfgclone.pl
dbconfig ${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml <<-EOF1
${APPS_PASSWD}
EOF1
"

```

16. For the EBS database Oracle 12.2, add a **Run Bash Shell Command** operation to the Configure Clone hook to ensure that `sqlnet.ora` or `sqlnet_ifile.ora` specify a value

for `SQLNET.ALLOWED_LOGON_VERSION_SERVER` for a **high privileged user**, as shown in the script below.

This requirement is outlined in **Cloning Oracle E-Business Suite Release 12.2 with Rapid Clone (Doc ID 1383621.1)** found at <http://docs.oracle.com>.

17. When prompted, if your environment is on R12.TXK.C.Delta.13 or later, enter the password for the EBS_SYSTEM user, if your environment is on R12.TXK.C.Delta.12 or earlier, enter the password for the SYSTEM user, or if your source EBS setup is running on AD-TXK Delta 13 then you need to supply the EBS_SYSTEM schema password in the below second hook script.

This requirement is outlined in the FAQ: Oracle E-Business Suite and System Schema Migration (DocID 2758999.1) Using UTL_FILE_DIR or Database Directories for PL/SQL File I/O in Oracle E-Business Suite Releases 12.1 and 12.2 (DocID 2525754.1) found at <http://docs.oracle.com>.

Configure Clone Second Hook for a 19c Multi-Tenant High Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# Set the target UTL_FILE_DIR values in the Oracle 19c database in the below
# variable UTL_FILE_DIR_PATH separated by commas
# Example UTL_FILE_DIR_PATH=/u01/oracle/VIS/temp/VISPDB,\
# /u01/oracle/VIS/19.3.0/appsutil/outbound/UTL,/u01/tmp

# UTL_FILE_DIR_PATH=<provide utl_file_dir locations separated by commas>
# If your environment is on R12.TXK.C.Delta.13 or later, enter the password
# for the EBS_SYSTEM user. If your environment is on R12.TXK.C.Delta.12 or
# earlier, enter the password for the SYSTEM user.
SYSTEM_PASSWD="<Password of system DB user>"
DLPX_HOSTNAME=$(hostname -s)
TIMESTAMP=$(date +%d-%m-%Y_%H-%M-%S)
CONTEXT_NAME=${DELPHIX_PDB_NAME}_${DLPX_HOSTNAME}
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
DIR_OBJ_PATH="${ORACLE_HOME}/appsutil/outbound/${DELPHIX_PDB_NAME}_${
DLPX_HOSTNAME}"

```

```

checkExists() {
if [[ ! -f $1 ]]; then
    echo "$1 does not exist."
    exit 1
fi
}
checkExists "${ORACLE_HOME}/${CONTEXT_NAME}.env"
checkExists "${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl"
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
echo "${APPS_PASSWD}" | perl "${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl"
    -contextfile="${ORACLE_HOME}/appsutil/${DELPHIX_PDB_NAME}_$
{DLPX_HOSTNAME}.xml" -oraclehome="${ORACLE_HOME}" -outdir="${ORACLE_HOME}/
appsutil/log" -mode=getUtlFileDir
checkExists "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt"
cp "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt" "${ORACLE_HOME}/dbs/
${DELPHIX_PDB_NAME}_utlfiledir.txt${TIMESTAMP}"
cat /dev/null > "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt"
if [ -z "$UTL_FILE_DIR_PATH" ]; then
    mkdir -p "${ORACLE_HOME}/appsutil/outbound/${DELPHIX_PDB_NAME}_$
{DLPX_HOSTNAME}"
    echo "${ORACLE_HOME}/appsutil/outbound/${DELPHIX_PDB_NAME}_$
{DLPX_HOSTNAME}" > "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt"
    mkdir -p "${ORACLE_BASE}/temp/${DELPHIX_PDB_NAME}"
    echo "${ORACLE_BASE}/temp/${DELPHIX_PDB_NAME}" >> "${ORACLE_HOME}/dbs/$
{DELPHIX_PDB_NAME}_utlfiledir.txt"
else
    echo "$UTL_FILE_DIR_PATH" | awk -F, '{ for (i = 1; i < NF+1; ++i ) print $i
>> "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt" }'
fi
{ echo "${APPS_PASSWD}"; echo "${SYSTEM_PASSWD}"; } | perl "${ORACLE_HOME}/
appsutil/bin/txkCfgUtlfileDir.pl" -contextfile="${ORACLE_HOME}/appsutil/$
{DELPHIX_PDB_NAME}_$ {DLPX_HOSTNAME}.xml" -oraclehome="${ORACLE_HOME}" -outdir="
${ORACLE_HOME}/appsutil/log" -mode=setUtlFileDir
{ echo "${APPS_PASSWD}"; echo "${SYSTEM_PASSWD}"; echo "${DIR_OBJ_PATH}"; } |
perl "${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl" -contextfile="$
{ORACLE_HOME}/appsutil/${DELPHIX_PDB_NAME}_$ {DLPX_HOSTNAME}.xml" -oraclehome="$
{ORACLE_HOME}" -outdir="${ORACLE_HOME}/appsutil/log" -mode=createDirObject
echo "${APPS_PASSWD}" | perl "${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl"
    -contextfile="${ORACLE_HOME}/appsutil/${DELPHIX_PDB_NAME}_$
{DLPX_HOSTNAME}.xml" -oraclehome="${ORACLE_HOME}" -outdir="${ORACLE_HOME}/
appsutil/log" -mode=syncUtlFileDir -skipautoconfig=yes

```

18. Configure Clone hook to ensure that `sqlnet.ora` or `sqlnet_ifile.ora` specify a value for `SQLNET.ALLOWED_LOGON_VERSION_SERVER` for a **low privileged user**, as shown in the script below.

When prompted, if your environment is on R12.TXK.C.Delta.13 or later, enter the password for the EBS_SYSTEM user, if your environment is on R12.TXK.C.Delta.12 or earlier, enter the password for the SYSTEM user, or if your source EBS setup is running on AD-TXK Delta 13 then you need to supply the EBS_SYSTEM schema password in the below second hook script.

This requirement is outlined in the FAQ: Oracle E-Business Suite and System Schema Migration (DocID 2758999.1) Using UTL_FILE_DIR or Database Directories for PL/SQL File I/O in Oracle E-Business Suite Releases 12.1 and 12.2 (DocID 2525754.1) found at <http://docs.oracle.com>.

Configure Clone Second Hook for a 19c Multi-Tenant Low Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# Set the target UTL_FILE_DIR values in the Oracle 19c database
# Example UTL_FILE_DIR_PATH=/u01/oracle/VIS/temp/VISPDB,\
# /u01/oracle/VIS/19.3.0/appsutil/outbound/UTL

# UTL_FILE_DIR_PATH=<provide utl_file_dir locations separated by commas>
# If your environment is on R12.TXK.C.Delta.13 or later, enter the password
# for the EBS_SYSTEM user. If your environment is on R12.TXK.C.Delta.12 or
# earlier, enter the password for the SYSTEM user.
SYSTEM_PASSWD="<Password of system DB user>"
DLPX_DB_EXEC_SCRIPT="<Remote BIN location till dlp_x_db_exec script>"
DLPX_PRIV_USER=oravis
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
CONTEXT_NAME=${DELPHIX_PDB_NAME}_$(hostname -s)
DIR_OBJ_PATH="${ORACLE_HOME}/appsutil/outbound/${CONTEXT_NAME}"
TIMESTAMP=$(date +%d-%m-%Y_%H-%M-%S)
checkExists() {
if [[ ! -f $1 ]]; then
    echo "$1 does not exist."
    exit 1
fi
}
checkExists "${ORACLE_HOME}/${CONTEXT_NAME}.env"
checkExists "${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl"
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${
CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl
-contextfile=${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml -oraclehome=${
ORACLE_HOME} -outdir=${ORACLE_HOME}/appsutil/log -mode=getUtlfileDir <<-EOF
${APPS_PASSWD}
EOF
"

checkExists "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt"
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "cp ${ORACLE_HOME}/dbs/${
DELPHIX_PDB_NAME}_utlfiledir.txt" "${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}
_utlfiledir.txt${TIMESTAMP}"
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "echo "" > \"${ORACLE_HOME}/
dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt\""

# Initialise for shell check to succeed
i=0
if [ -z "$UTL_FILE_DIR_PATH" ]; then

```

```

    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "mkdir -p ${ORACLE_HOME}/
appsutil/outbound/${CONTEXT_NAME}"
    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "echo ${ORACLE_HOME}/
appsutil/outbound/${CONTEXT_NAME} > ${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}
_utlfiledir.txt"
    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "mkdir -p ${ORACLE_BASE}/
temp/${DELPHIX_PDB_NAME}"
    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "echo ${ORACLE_BASE}/temp/${
DELPHIX_PDB_NAME} >> ${ORACLE_HOME}/dbs/${DELPHIX_PDB_NAME}_utlfiledir.txt"
else
    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "echo $UTL_FILE_DIR_PATH |
awk -F, '{ for (i = 1; i < NF+1; ++i ) print $i }' > ${ORACLE_HOME}/dbs/${
DELPHIX_PDB_NAME}_utlfiledir.txt"
fi

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl
-contextfile=${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml -oraclehome=$
{ORACLE_HOME} -outdir=${ORACLE_HOME}/appsutil/log -mode=setUtlfileDir <<-EOF1
${APPS_PASSWD}
${SYSTEM_PASSWD}
EOF1
"

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl
-contextfile=${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml -oraclehome=$
{ORACLE_HOME} -outdir=${ORACLE_HOME}/appsutil/log -mode=createDirObject <<-EOF2
${APPS_PASSWD}
${SYSTEM_PASSWD}
${DIR_OBJ_PATH}
EOF2
"

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/bin/txkCfgUtlfileDir.pl
-contextfile=${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml -oraclehome=$
{ORACLE_HOME} -outdir=${ORACLE_HOME}/appsutil/log -mode=syncUtlfileDir
-skipautoconfig=yes <<-EOF3
${APPS_PASSWD}
EOF3
"

```

19. Configure Clone hook to ensure that the ebs_<PDB_SID> listener services are not disappeared at the stop or start of the VDB, as shown in the script below.

Third Configure Clone Hook for a 19c Multi-Tenant High Privileged and Low Privileged Users

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
. "${ORACLE_HOME}/${ORACLE_SID}_"$(hostname -s).env";

```

```
sqlplus "/" as sysdba" <<-EOF
alter pluggable database ${DELPHIX_PDB_NAME} close;
alter pluggable database ${DELPHIX_PDB_NAME} open read write
services=all;
alter pluggable database ${DELPHIX_PDB_NAME} save state;
EOF
```

20. Set up a Pre-Snapshot hook **Run Bash Shell Command** operation to run any pre-clone steps necessary and specific to your EBS database for a **high privileged user**, as shown in the script below.

Normally, these steps will include running Oracle's `adpreclone` tool.

Pre-Snapshot Hook for a 19c Multi-Tenant High Privileged User

```
#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# As grep has to check all the matches in making decision on L58
# shellcheck disable=SC2143
# NOTE: Ensure the below environment variables will be set up correctly by
# the shell. If not, hardcode or generate the values below.
CONTEXT_NAME=${DELPHIX_PDB_NAME}_$(hostname -s)
SOURCE_APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
TARGET_APPS_PASSWD="<Target apps schema password>"
timeout=3600
waittime=0
ORACLE_USER="oravis"

# Query active PDB services
. "${ORACLE_HOME}/${ORACLE_SID}_$(hostname -s).env";
check_value=$(sqlplus -s "/" as sysdba" <<EOF
SELECT NAME FROM v\$active_services;
exit;
EOF
)

checkService() {
if ! echo "$1" | grep -q "$2"; then
. "${ORACLE_HOME}/${ORACLE_SID}_$(hostname -s).env";
sqlplus -s "/" as sysdba" <<EOF
alter session set container="${DELPHIX_PDB_NAME}";
BEGIN DBMS_SERVICE.START_SERVICE( service_name => '$2');
end;
/
EOF
fi
}

# Check for ebs_PDB service
```

```

checkService "$check_value" "ebs_${DELPHIX_PDB_NAME}"
# Check for PDB_ebs_patch service
checkService "$check_value" "${DELPHIX_PDB_NAME}_ebs_patch"

. "${ORACLE_HOME}/${CONTEXT_NAME}.env";

testAppsPassword() {
    local passwordInQuestion=$1
    ERROR=$(sqlplus "apps/${passwordInQuestion}@${DELPHIX_PDB_NAME}" <<<
"exit;")
    grep ORA-01017 <<< "${ERROR}" >/dev/null && return 1
    return 0
}

testAppsPassword "${SOURCE_APPS_PASSWD}"
testResult=$?

if [[ ${testResult} == 0 ]]; then
    APPS_PASSWD="${SOURCE_APPS_PASSWD}"
else
    APPS_PASSWD="${TARGET_APPS_PASSWD}"
fi

while [[ -f "${ORACLE_HOME}/.delphix_adpreclone.lck" || $(pgrep -au "${ORACLE_USER}" | grep "^${ORACLE_HOME}" | grep -v grep | grep "adpreclone") ]]; do
    if [[ $waittime -gt $timeout ]]; then
        echo "Another adpreclone process is still running from last 60 mins.Delphix cannot proceed until it is complete. Exiting Now."
        exit 1
    fi
    echo " Another adpreclone process is running. waiting for the process to complete before starting adpreclone of the database...."
    (( timeout += 10 ))
    sleep 10
done
echo "No other adpreclone process is running on database, proceeding ...."
"${ORACLE_HOME}"/perl/bin/perl "${ORACLE_HOME}/appsutil/scripts/${CONTEXT_NAME}/adpreclone.pl" database <<-EOF
${APPS_PASSWD}
EOF

```

Pre-Snapshot hook **Run Bash Shell Command** operation to run any pre-clone steps necessary and specific to your EBS database for a **low privileged user**, as shown in the script below.

i `${ORACLE_SID}_source_apps_file.txt` will be created automatically inside the `/var/tmp/` location of the target database host once the DBTechStack provisioning is completed.

Pre-Snapshot Hook for a 19c Multi-Tenant Low Privileged User

```
#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# As grep has to check all the matches in making decision on L58
# shellcheck disable=SC2143
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.
DLPX_DB_EXEC_SCRIPT="<Remote BIN location till dlp_x_db_exec script>"
DLPX_PRIV_USER=oravis
CONTEXT_NAME=${DELPHIX_PDB_NAME}_$(hostname -s)
SOURCE_APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
TARGET_APPS_PASSWD="<target apps schema password>"
timeout=3600
waittime=0

# Query active PDB services
. "${ORACLE_HOME}/${ORACLE_SID}_$(hostname -s).env";
check_value=$(sqlplus -s "/ as sysdba" <<EOF
SELECT NAME FROM v\$active_services;
exit;
EOF
)

checkService() {
if ! echo "$1" | grep -q "$2"; then
. "${ORACLE_HOME}/${ORACLE_SID}_$(hostname -s).env";
sqlplus -s "/ as sysdba" <<EOF
alter session set container="${DELPHIX_PDB_NAME}";
BEGIN DBMS_SERVICE.START_SERVICE( service_name => '$2');
end;
/
EOF
fi
}

# Check for ebs_PDB service
checkService "$check_value" "ebs_${DELPHIX_PDB_NAME}"
# Check for PDB_ebs_patch service
checkService "$check_value" "${DELPHIX_PDB_NAME}_ebs_patch"

. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
testAppsPassword() {
    local passwordInQuestion=$1
    ERROR=$( "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${
CONTEXT_NAME}.env; sqlplus apps/\ "${passwordInQuestion}"@"${DELPHIX_PDB_NAME} <<<
\ "exit;\ " )
    grep ORA-01017 <<< "${ERROR}" >/dev/null && return 1

```

```

    return 0
}
testAppsPassword "${SOURCE_APPS_PASSWD}"
testResult=$?

if [[ ${testResult} == 0 ]]; then
    APPS_PASSWD=${SOURCE_APPS_PASSWD}
else
    APPS_PASSWD=${TARGET_APPS_PASSWD}
fi
while [[ -f "${ORACLE_HOME}/delphix_adpreclone.lck" || $(pgrep -au "$
{DLPX_PRIV_USER}" | grep "^${ORACLE_HOME}" | grep -v grep | grep "adpreclone") ]];
do
    if [[ $waittime -gt $timeout ]]; then
        echo "Another adpreclone process is still running from last 60 mins.Delphix
cannot proceed until it is complete. Exiting Now."
        exit 1
    fi
    echo " Another adpreclone process is running. waiting for the process to complete
before starting adpreclone of the database...."
    (( timeout += 10 ))
    sleep 10
done
echo "No other adpreclone process is running on database, proceeding ...."
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${CONTEXT_NAME}.env;
${ORACLE_HOME}/perl/bin/perl ${ORACLE_HOME}/appsutil/scripts/${CONTEXT_NAME}/
adpreclone.pl database <<-EOF
${APPS_PASSWD}
EOF
"

```

Provisioning the Oracle 12C Database

1. Refresh the Target DB environment from **Manage > Environments** so that the Listener brought up during dbTechStack provision gets discovered and the Installation Home to be listed in dropdown during Database provision. If Target DB environment is not refreshed, you will see `This environment has no compatible Oracle Installation Homes` in the Installation Home dropdown in the **Database provision**.
2. Provision the EBS database to the target dbTier environment by following the steps outlined in [Provisioning an Oracle VDB](#). For Oracle EBS database for 19c MT, provision the database by following the steps outlined in [Provisioning an Oracle Virtual Pluggable Database](#).

Snapshot conflicts

When Snapshot is running against the dbTechStack, database, or appsTier, the Delphix Engine also executes pre-clone logic to ensure the latest configuration is staged in the captured snapshots. Unfortunately, if multiple Snapshots are running against the same EBS instance concurrently, this pre-clone logic may fail and produce bad snapshots. To avoid SnapSync conflicts, spread out your SnapSync policies for an EBS instance by one hour or more.

3. Select the correct **Installation Home**. This should be the virtual dbTechStack you just added to the Delphix Engine.
4. Select an Environment User.

5. This user should be the oracle user-outlined in Preparing Target EBS R12.2 Environments for Provisioning.
6. Select a Target Group for the VDB.
7. In the VDB **Configuration** section, provide Target DB/CDB SID name provided during dbTechStack provisioning in the **Oracle Database Name** and **Oracle SID** field. Provide any unique user-defined name in **VDB Name** and **Oracle Database Unique Name** and click **Next**.
8. Click **Advanced**.
9. Select the correct **Oracle Node Listeners** value.
This should be the listener corresponding to the virtual dbTechStack you just added to the Delphix Engine.
10. Add the EBS R12.2 dbTier environment file as a **Custom Environment Variables** entry.
11. This file can be specified as an Environment File with Path Parameters of `$_ORACLE_HOME/
<CONTEXT_NAME>.env` .
Replace `<CONTEXT_NAME>` with the virtual EBS instance's context name. The Delphix Engine will expand the `$_ORACLE_HOME` variable at runtime.
For more information, see [Customizing Oracle VDB Environment Variables](#).
12. Add a **Run Bash Shell Command** operation to the Configure Clone hook to ensure that `adcfgclone` is run against the newly provisioned database for a **high privileged user**, as shown in the script below.

Configure Clone First Hook for a High Privileged User

```
#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# NOTE: Ensure the below environment variables will be set up correctly by
# the shell. If not, hardcode or generate the values below.
HOSTNAME=$(uname -n | cut -d '.' -f1)
CONTEXT_NAME=${ORACLE_SID}_${HOSTNAME}
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")

. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
# Check for local_listener parameter is set for SID, otherwise set it
# appropriately
check_value=$(sqlplus -s "/" as sysdba" <<EOF
show parameter local_listener;
EOF
)
local_listener=$(echo "$check_value" | awk '{print $11}')
value=("${local_listener//:/ }")
host="${value[0]}"
port="${value[1]}"
curr_port=$(grep PORT < "${TNS_ADMIN}/listener.ora" | awk '{print $9}' | sed
's/)//g')

if [[ $port != "$curr_port" || $host != "${HOSTNAME}" ]]; then
sqlplus -s "/" as sysdba" <<EOF
alter system set local_listener='${HOSTNAME}:${curr_port}';
alter system register;
EOF
```

```

fi

sqlplus "/" as sysdba" <<EOF
@${ORACLE_HOME}/appsutil/install/${CONTEXT_NAME}/adupdlib.sql so
EOF

perl "${ORACLE_HOME}/appsutil/clone/bin/adcfgclone.pl" dbconfig "$
${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml" <<EOF
${APPS_PASSWD}
EOF

```

13. Configure Clone hook to ensure that `adcfgclone` is run against the newly provisioned database for a **low privileged user**, as shown in the script below.

 Customers have to identify the `DLPX_DB_EXEC_SCRIPT` value from their remote environment, this would always be inside the toolkit directory provided while adding the environment on Delphix Engine.

Configure Clone First Hook for a Low Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.

DLPX_DB_EXEC_SCRIPT="<Remote BIN location for dlp_x_db_exec script>"
DLPX_PRIV_USER=oravis
APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
HOSTNAME=$(uname -n | cut -d '.' -f1)
CONTEXT_NAME=${ORACLE_SID}_${HOSTNAME}

. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
# Check for local_listener parameter is set for SID, otherwise set it
# appropriately
check_value=$(sqlplus -s "/" as sysdba" <<EOF
show parameter local_listener;
EOF
)
local_listener=$(echo "$check_value" | awk '{print $11}')
value=("${local_listener//:/ }")
host="${value[0]}"
port="${value[1]}"
curr_port=$(grep PORT < "${TNS_ADMIN}/listener.ora" | awk '{print $9}' | sed
's/)//g')

if [[ $port != "$curr_port" || $host != "${HOSTNAME}" ]]; then

```

```

sqlplus -s "/ as sysdba" <<EOF
alter system set local_listener='${HOSTNAME}:${curr_port}';
alter system register;
EOF
fi

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{CONTEXT_NAME}.env; sqlplus \"/ as sysdba\" <<-EOF
@${ORACLE_HOME}/appsutil/install/${CONTEXT_NAME}/adupdlib.sql so
EOF
"

"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/$
{CONTEXT_NAME}.env; perl ${ORACLE_HOME}/appsutil/clone/bin/adcfgclone.pl
dbconfig ${ORACLE_HOME}/appsutil/${CONTEXT_NAME}.xml <<-EOF1
${APPS_PASSWD}
EOF1
"

```

14. For the EBS database Oracle 12.2, add a **Run Bash Shell Command** operation to the Configure Clone hook to ensure that `sqlnet.ora` or `sqlnet_ifile.ora` specify a value for `SQLNET.ALLOWED_LOGON_VERSION_SERVER` for a **high privileged user**, as shown in the script below.

This requirement is outlined in **Cloning Oracle E-Business Suite Release 12.2 with Rapid Clone (Doc ID 1383621.1)** found at <http://docs.oracle.com>.

Configure Clone Second Hook for a High Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# NOTE: Ensure the below environment variables will be set up correctly by
# the shell. If not, hardcode or generate the values below.
# If you are using sqlnet_ifile.ora, change the script below to reflect
# sqlnet_ifile.ora Initialize dummy variable `parameter` to avoid shellcheck
# to fail
parameter=""
CONTEXT_NAME=${ORACLE_SID}_(uname -n | cut -d '.' -f1)
TNS_ADMIN=${ORACLE_HOME}/network/admin/${CONTEXT_NAME}
check_value=$(sqlplus -s "/ as sysdba" <<EOF
set head off termout off feedback off wrap off
select DISPLAY_VALUE from v\\$parameter where NAME='sec_case_sensitive_logon';
EOF
)

if [[ ${check_value} = "FALSE" ]]; then
    sed 's/^SQLNET.ALLOWED_LOGON_VERSION_SERVER=.*/'
SQLNET.ALLOWED_LOGON_VERSION_SERVER=8/g' "${TNS_ADMIN}/sqlnet.ora" > /var/tmp/
temp.ora && mv /var/tmp/temp.ora "${TNS_ADMIN}/sqlnet.ora"

```

```

elif [[ ${check_value} = "TRUE" ]]; then
    sed 's/^SQLNET.ALLOWED_LOGON_VERSION_SERVER=.*\nSQLNET.ALLOWED_LOGON_VERSION_SERVER=10/g' "${TNS_ADMIN}/sqlnet.ora" > /var/tmp/temp.ora && mv /var/tmp/temp.ora "${TNS_ADMIN}/sqlnet.ora"
else
    echo "sec_case_sensitive_logon parameter is not set in the database. So the sqlnet.ora has not been updated."
fi

```

15. When prompted, if your environment is on R12.TXK.C.Delta.13 or later, enter the password for the EBS_SYSTEM user, if your environment is on R12.TXK.C.Delta.12 or earlier, enter the password for the SYSTEM user, or if your source EBS setup is running on AD-TXK Delta 13 then you need to supply the EBS_SYSTEM schema password in the below second hook script. This requirement is outlined in the FAQ: Oracle E-Business Suite and System Schema Migration (DocID 2758999.1) Using UTL_FILE_DIR or Database Directories for PL/SQL File I/O in Oracle E-Business Suite Releases 12.1 and 12.2 (DocID 2525754.1) found at <http://docs.oracle.com>.
16. Configure Clone hook to ensure that `sqlnet.ora` or `sqlnet_ifile.ora` specify a value for `SQLNET.ALLOWED_LOGON_VERSION_SERVER` for a **low privileged user**, as shown in the script below.

Configure Clone Second Hook for a Low Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.
# If you are using sqlnet_ifile.ora, change the script below to reflect
# sqlnet_ifile.ora.
# Initialize dummy variable `parameter` to avoid shellcheck to fail
parameter=""
DLPX_DB_EXEC_SCRIPT="<Remote BIN location for dlp_x_db_exec script>"
DLPX_PRIV_USER=oravis
CONTEXT_NAME=${ORACLE_SID}_${(uname -n | cut -d '.' -f1)}
TNS_ADMIN=${ORACLE_HOME}/network/admin/${CONTEXT_NAME}

check_value=$(("${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/
${CONTEXT_NAME}.env; sqlplus \"/ as sysdba\" <<-EOF
set head off termout off feedback off wrap off
select DISPLAY_VALUE from v\\$parameter where NAME='sec_case_sensitive_logon';
EOF
")

if [[ ${check_value} = "FALSE" ]]; then
    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "sed 's/
^SQLNET.ALLOWED_LOGON_VERSION_SERVER=.*\nSQLNET.ALLOWED_LOGON_VERSION_SERVER=8/
g' \("${TNS_ADMIN}/sqlnet.ora\" > /var/tmp/temp.ora && mv /var/tmp/temp.ora \("${
{TNS_ADMIN}/sqlnet.ora\""
elif [[ ${check_value} = "TRUE" ]]; then

```

```

    "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" "sed 's/
^SQLNET.ALLOWED_LOGON_VERSION_SERVER=.*$/SQLNET.ALLOWED_LOGON_VERSION_SERVER=10/
g' \"${TNS_ADMIN}/sqlnet.ora\" > /var/tmp/temp.ora && mv /var/tmp/temp.ora \"${
TNS_ADMIN}/sqlnet.ora\""
else
    echo "sec_case_sensitive_logon parameter is not set in the database. So the
sqlnet.ora has not been updated."
fi

```

When prompted, if your environment is on R12.TXK.C.Delta.13 or later, enter the password for the EBS_SYSTEM user, if your environment is on R12.TXK.C.Delta.12 or earlier, enter the password for the SYSTEM user, or if your source EBS setup is running on AD-TXK Delta 13 then you need to supply the EBS_SYSTEM schema password in the below second hook script.

This requirement is outlined in the FAQ: Oracle E-Business Suite and System Schema Migration (DocID 2758999.1) Using UTL_FILE_DIR or Database Directories for PL/SQL File I/O in Oracle E-Business Suite Releases 12.1 and 12.2 (DocID 2525754.1) found at <http://docs.oracle.com>.

17. Configure Clone hook to ensure that the ebs_<PDB_SID> listener services are not disappeared at the stop or start of the VDB, as shown in the script below.
18. Set up a Pre-Snapshot hook **Run Bash Shell Command** operation to run any pre-clone steps necessary and specific to your EBS database for a **high privileged user**, as shown in the script below.

Normally, these steps will include running Oracle's `adpreclone` tool.

Pre-Snapshot Hook for a High Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# As grep has to check all the matches in making decision on L58
# shellcheck disable=SC2143
# NOTE: Ensure the below environment variables will be set up correctly by
# the shell. If not, hardcode or generate the values below.
HOSTNAME=$(uname -n | cut -d '.' -f1)
CONTEXT_NAME=${ORACLE_SID}_${HOSTNAME}
SOURCE_APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
TARGET_APPS_PASSWD="<new apps password of the target DB node>"
ORACLE_USER="oravis"

timeout=3600
waittime=0
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";

testAppsPassword() {
    local passwordInQuestion=$1
    ERROR=$(sqlplus "apps/${passwordInQuestion}" <<< "exit;")

    grep ORA-01017 <<< "${ERROR}" >/dev/null && return 1
    return 0
}

```

```

}

testAppsPassword "${SOURCE_APPS_PASSWD}"
testResult=$?

if [[ ${testResult} == 0 ]]; then
    APPS_PASSWD=${SOURCE_APPS_PASSWD}
else
    APPS_PASSWD=${TARGET_APPS_PASSWD}
fi

# Query active SID services
check_value=$(sqlplus -s "/ as sysdba" <<EOF
SELECT NAME FROM v\$active_services;
exit;
EOF
)
checkService() {
if ! echo "$1" | grep -q "$2"; then
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
sqlplus -s "/ as sysdba" <<EOF
BEGIN DBMS_SERVICE.START_SERVICE( service_name => '$2');
end;
/
EOF
fi
}
# Check for SID_ebs_patch or ebs_patch service based on the AD-TXK codelevel
# version. In AD-TXK Delta 8 and earlier, the service name for connections
# to the patch edition of the database was ebs_patch. In AD-TXK Delta 9, the
# service name to connect to the patch edition has been changed to
# <instance_name>_ebs_patch
patch_level=$(sqlplus -s "apps/${APPS_PASSWD}" <<EOF
select PATCH_LEVEL from apps.fnd_product_installations where PATCH_LEVEL like
'R12.AD.%';
EOF
)
patch_service="${ORACLE_SID}_ebs_patch"
if [ "${patch_level: -1}" -le 8 ]; then
    patch_service="ebs_patch"
fi
checkService "$check_value" "${patch_service}"

while [[ -f "${ORACLE_HOME}/.delphix_adpreclone.lck" || $(pgrep -au "$ORACLE_USER" | grep "^${ORACLE_HOME}" | grep -v grep | grep "adpreclone") ]]
; do
    if [[ $waittime -gt $timeout ]]; then
        echo "Another adpreclone process is still running from last 60
mins.Delphix cannot proceed until it is complete. Exiting Now."
        exit 1
    fi
fi

```

```

    echo " Another adpreclone process is running. waiting for the process to
complete before starting adpreclone of the database...."
    (( timeout += 10 ))
    sleep 10
done
echo "No other adpreclone process is running on database, proceeding ...."
"${ORACLE_HOME}"/perl/bin/perl "${ORACLE_HOME}/appsutil/scripts/$
{CONTEXT_NAME}/adpreclone.pl" database <<-EOF
${APPS_PASSWD}
EOF

```

19. Pre-Snapshot hook **Run Bash Shell Command** operation to run any pre-clone steps necessary and specific to your EBS database for a **low privileged user**, as shown in the script below.

Pre-Snapshot Hook for a Low Privileged User

```

#!/usr/bin/env bash
#
# Copyright (c) 2022 by Delphix. All rights reserved.
#
# shellcheck source=/dev/null
# NOTE: Ensure the below environment variables will be set up correctly by the
# shell. If not, hardcode or generate the values below.
# As grep has to check all the matches in making decision on L58
# shellcheck disable=SC2143
DLPX_DB_EXEC_SCRIPT=""
DLPX_PRIV_USER=oravis
HOSTNAME=$(uname -n | cut -d '.' -f1)
CONTEXT_NAME=${ORACLE_SID}_${HOSTNAME}
SOURCE_APPS_PASSWD=$(cat "/var/tmp/${ORACLE_SID}_source_apps_file.txt")
TARGET_APPS_PASSWD=""

timeout=3600
waittime=0

testAppsPassword() {
    local passwordInQuestion=$1
    ERROR=$( "${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${
CONTEXT_NAME}.env; sqlplus apps/\ "${passwordInQuestion}" <<< \ "exit;\")

    grep ORA-01017 <<< "${ERROR}" >/dev/null && return 1
    return 0
}

testAppsPassword "${SOURCE_APPS_PASSWD}"
testResult=?

if [[ ${testResult} == 0 ]]; then
    APPS_PASSWD=${SOURCE_APPS_PASSWD}
else
    APPS_PASSWD=${TARGET_APPS_PASSWD}

```

```

fi

# Query active SID services
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
check_value=$(sqlplus -s "/" as sysdba" <<EOF
SELECT NAME FROM v\$active_services;
exit;
EOF
)
checkService() {
if ! echo "$1" | grep -q "$2"; then
. "${ORACLE_HOME}/${CONTEXT_NAME}.env";
sqlplus -s "/" as sysdba" <<EOF
BEGIN DBMS_SERVICE.START_SERVICE( service_name => '$2');
end;
/
EOF
fi
}
# Check for SID_ebs_patch or ebs_patch service based on the AD-TXK codelevel
# version. In AD-TXK Delta 8 and earlier, the service name for connections
# to the patch edition of the database was ebs_patch. In AD-TXK Delta 9, the
# service name to connect to the patch edition has been changed to
# <instance_name>_ebs_patch
patch_level=$(sqlplus -s "apps/${APPS_PASSWD}" <<EOF
select PATCH_LEVEL from apps.fnd_product_installations where PATCH_LEVEL like
'R12.AD.%';
EOF
)
patch_service="${ORACLE_SID}_ebs_patch"
if [ "${patch_level: -1}" -le 8 ]; then
    patch_service="ebs_patch"
fi
checkService "$check_value" "${patch_service}"

while [[ -f "${ORACLE_HOME}/.delphix_adpreclone.lck" || $(pgrep -au "$
{DLPX_PRIV_USER}" | grep "^${ORACLE_HOME}" | grep -v grep | grep "adpreclone")
]] ; do
    if [[ $waittime -gt $timeout ]]; then
        echo "Another adpreclone process is still running from last 60
mins.Delphix cannot proceed until it is complete. Exiting Now."
        exit 1
    fi
    echo " Another adpreclone process is running. waiting for the process to
complete before starting adpreclone of the database...."
    (( timeout += 10 ))
    sleep 10
done
echo "No other adpreclone process is running on database, proceeding ...."
"${DLPX_DB_EXEC_SCRIPT}" -u"${DLPX_PRIV_USER}" ". ${ORACLE_HOME}/${
CONTEXT_NAME}.env; ${ORACLE_HOME}/perl/bin/perl ${ORACLE_HOME}/appsutil/
scripts/${CONTEXT_NAME}/adpreclone.pl database <<-EOF

```

```

${APPS_PASSWD}
EOF
"

```

i `_${ORACLE_SID}_source_apps_file.txt` will be created automatically inside the `/var/tmp/` location of the target database host once the DBTechStack provisioning is completed.

Provisioning the EBS appsTier

1. Login to the **Delphix Management** application using **Admin** credentials.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **appsTier dSource**.
5. Select a dSource **snapshot**.

All snapshots will have staged configuration prepared by `adpreclone.pl` and any hook operations placed on the dSource.

6. Click **Provision**.
The **Provision VDB** wizard will open.
7. Select an **Environment**.
This environment will host the virtual appsTier and be used to execute hook operations specified in a few steps. This environment will also run the WebLogic Admin server (Web Administration service) for the virtual appsTier.
If you are provisioning a multi-node appsTier, you will be able to specify additional environments to host the virtual appsTier in a few steps.
8. Select E-Business Suite R12.2 appsTier from the **Installation** dropdown.
9. Select an Environment User.

This user should be the `app_lmgr` user-outlined in [Preparing Target EBS R12.2 Environments for Provisioning](#).

10. Enter a **Mount Path** for the virtual appsTier VDB.
If you are provisioning a multi-node appsTier, this mount path will be used across all target environments.
11. Enter the **EBS-specific parameters** for the virtual appsTier. A subset of these parameters are discussed in more detail below.
 - a. **Privileged OS Account (Optional)** field should contain high privileged user when low privileged user is being used for provisioning.
 - b. The **Target APPS Password** is the new apps password that is required to configure and manage the virtual appsTier.
This password is encrypted when stored within the Delphix Engine and is available as an environment variable to the `adcfgclone`, `adstrtal`, and `adstpall` processes.
 - c. The **Weblogic AdminServer Password** is the new WebLogic password required to configure the virtual appsTier. A password change will be performed if this password does not match that of the source EBS instance. This password is encrypted when stored within the Delphix Engine and is available as an environment variable to the `adcfgclone`, `adstrtal`, and `adstpall` processes.
 - d. The **SYSTEM Password** is required to configure the virtual appsTier with the new apps password.
This password is not required if the Target Apps Password is the same as the source.
 - e. Provide the Target Application hostname.
 - f. The **Target DB/PDB SID** is the new database SID (PDB SID in case of 19c) that is required to configure and manage the virtual appsTier.
 - g. Provide the Target DB server node hostname.
 - h. Provide **Target DB Domain Name**, this will be provided to the `adcfgclone`.
 - i. Delphix recommends specifying an **Instance Home Directory** under the **Mount Path** so that instance-specific EBS files live on Delphix-provided storage.

For example, if the provided **Mount Path** is `/u01/oracle/VIS`, then providing an **Instance Home Directory** of `/u01/oracle/VIS` would allow EBS to generate virtual application `INST_TOP` in `/u01/oracle/VIS/fs1/inst/apps/<CONTEXT_NAME >` and `/u01/oracle/VIS/fs2/inst/apps/<CONTEXT_NAME >`.

- i. If you are provisioning a single-node appsTier, this recommendation is OPTIONAL; putting instance-specific EBS files on Delphix-provided storage merely eases the administration of the virtual EBS instance.
- ii. If you are provisioning a multi-node appsTier, this recommendation is REQUIRED; the Delphix Engine's automation requires that all nodes in the appsTier have access to instance-specific files via Delphix-provided storage.
- j. Ensure that the **Target Application Hostname** and **Target DB Server Node** values are the short hostnames, not the fully-qualified hostnames.
- k. Select **Services**, by default all the **Root**, **Web Entry Point**, **Web Application** and **Batch Processing** would be checked. In case you are using multinode environment and want to start a few services on a secondary node then this field is going to be very useful for you.
 - l. Provide **DISPLAY Variable**, the default value is hostname:0.0
- m. Provide **Target Run Edition Port Pool**, the value of this field should be 0 to 99. This should match with the port pool value given during the DBTechStack provisioning.
- n. Provide **Target Patch Edition Port Pool**, the value of this field should be 0 to 99. This value should be Target Run Edition Port Pool plus 1. For example, if Run Edition Port Pool is 9 then the value should be 10.
- o. The **EBS AppsTier Timeout** is required to terminate all the long-running appsTier processes which have exceeded the timeout value when stopping the applications as part of a refresh. This timeout will be calculated in minutes. For example, if set to 30, then we run `adstpall` at the start of the refresh, and if after 30 minutes the application has not stopped, then the processes will be terminated to allow the refresh to continue.
- p. The **EBS AppsTier Provision/Refresh Timeout** is required to terminate configure and cloning via `adcfgclone` process which have exceeded the timeout as part of a provision/refresh. This timeout will be calculated in hours and the default value will be 8 hours. For example, if set to 9, then plugin will detect a timeout after 9 hours instead of getting in a hung state and UI error for the same will be prompted.
- q. The **Java Color Scheme** is an **optional** parameter and is used to change the JAVA color for provisioned Oracle EBS JAVA based form interface. Please choose any of the following colors (Swan, Blue, Khaki, Olive, Purple, Red, Teal, Titanium) to set Java Color Scheme for EBS appsTier java-based forms interface.
- r. Provide **Target System Proxy Hostname (Optional)**, This is an optional parameter that will be used if `adcfgclone` requires it.
- s. Provide **Target System Proxy Port (Optional)**, This is an optional parameter that will be used if `adcfgclone` requires it.
- t. Enable the **Cleanup Before Provision** option if you want to permit the Delphix Engine to automatically cleanup stale EBS configuration during a refresh. This option is recommended, but only available if your Oracle Home is patched with Oracle Universal Installer (OUI) version 10.2 or above.
 - i. With this option enabled, the Delphix Engine will inspect the target environment's `oraInventory` prior to refreshing this virtual appsTier. If any Oracle Homes are already registered within the specified **Mount Path**, the Delphix Engine will detach them from the inventory prior to running `adcfgclone`. These homes must be detached prior to running post-clone configuration. If they are not detached, `adcfgclone` will fail, citing conflicting `oraInventory` entries as an issue. The Delphix Engine will also remove any conflicting

- INST_TOP directories left on the environment. Non-conflicting INST_TOP directories will not be modified.
- ii. Without this option enabled, Oracle Homes or INST_TOP directories that conflict with the specified **Mount Path** or desired INST_TOP location will be reported in errors instead of automatically cleaned up. For refresh to succeed, you must manually detach conflicting Oracle Homes and manually remove conflicting INST_TOP directories prior to refresh.
 - u. Enable the **Start Services After Provision** option if you want to permit the Delphix Engine to automatically start the services for EBS appsTier after provisioning.
 - i. With this option enabled, the Delphix Engine will start all the services of EBS appsTier after provisioning.
 - ii. Without this option enabled, the Delphix Engine will **NOT** start the services of EBS appsTier after provisioning. This will allow customers to prevent EBS services from being started at the conclusion of provisioning or refreshing. That way, customers can perform post-clone processing automation using a configure-clone hook without having to stop services first.
 - v. If you are provisioning a multi-node appsTier, enter additional appsTier nodes as **Additional Nodes**.
 - i. Select the Environment for the secondary node.
 - ii. The **Environment User** for each node should be the `app\mgr` user-outlined in [Preparing Target EBS R12.2 Environments for Provisioning](#).
 - iii. Ensure that the **Hostname** value for each node is the short hostname, not the fully-qualified hostname.
 - iv. Provide DISPLAY Variable, the default value is hostname:0.0.
 - v. The **Mount Path** is not configurable for each node individually. The **Mount Path** provided for the primary environment will be used for each additional node.
12. Click **Next**.
 13. Enter a **VDB Name**.
 14. Select a **Target Group** for the VDB.
If necessary, click the **Plus** icon to add a new group.
 15. Select a **Snapshot Policy** for the VDB.
If necessary, click the **Plus** icon to create a new policy.

 When Snapshot is running against the dbTechStack, database, or appsTier, the Delphix Engine also executes pre-clone logic to ensure the latest configuration is staged in the captured snapshots. Unfortunately, if multiple Snapshots are running against the same EBS instance concurrently, this pre-clone logic may fail and produce bad snapshots. To avoid SnapSync conflicts, spread out your SnapSync policies for an EBS instance by one hour or more.

- Click **Next**.
16. Enter any **custom hook operations** that are needed to help correctly manage the virtual appsTier. The Configure Clone hook will be run after the `adcfgclone.pl` tool has both mounted and configured the appsTier.
All hook operations run against the environment specified for provision. For a multi-node appsTier, hook operations never run against additional nodes specified.
 17. Click **Next**.
 18. Click **Submit**.

 **dbTier Must Be Accessible During appsTier Provisioning**

Post-clone configuration will fail if the appsTier cannot connect to the database. Ensure the target dbTier is accessible to the appsTier during the provisioning process. Ensure both the virtual database and database listener are running.

When provisioning starts, you can review the progress of the job in the **Datasets** panel, or in the **Job History** panel of the **Dashboard**. When provisioning is complete, the appsTier VDB will be included in the group you designated and listed in the **Datasets** panel. If you select the appsTier VDB in the **Datasets** panel and click the **Configuration** tab, you can view information about the virtual files and its Data Management settings.

19. For tips on monitoring the progress of appsTier provisioning, see [Monitoring EBS R12.2 appsTier Provisioning Progress](#).

Once all three EBS virtual datasets have been provisioned successfully, your virtual EBS instance should be running and accessible.

Rewind Operation on Snapshot

Perform the rewind operation on the snapshot after Appstier provisioning in the following sequence.

1. Take a snapshot of the Database VDB after Appstier provision.
2. Stop the appsTier, DB, and dbTechStack one by one.
3. Rewind dbTechStack, DB, and appsTier one by one.

If the DB snapshot was taken before appsTier provisioning, rewinding the DB will cause loss of new data updated by appsTier during provisioning and the next operation of appsTier rewind will fail. So, to ensure that the Rewind operation of appsTier succeeds, take a snapshot of DB after Appstier provision.

Registering the EBS dbTechStack

Register the freshly-provisioned dbTechStack with the Delphix Engine.

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Environments**.
4. Select the **target dbTier environment**.
5. Click the **Databases** tab.
6. Click the **Plus** icon next to **Add Dataset Home**.
7. Select **Oracle** as your **Dataset Home Type**.
8. Enter an **Installation Home**.

This path should be the value of `$ORACLE_HOME` on your target dbTier; this path will live under the mount path specified when you provisioned the virtual dbTechStack.

Commonly, this path looks like `/u01/oracle/VIS/11.2.0`.

9. Click the **checkmark** to save your dataset home.
If necessary, scroll down the list of dataset homes to view and edit this dataset home.
 10. Refresh the dbTier environment.
Refreshing the environment will discover an EBS database listener and ensure it is available for selection when provisioning the EBS database.
1. Click the **Refresh** button in the top right-hand corner of the environment card.

Related Topics

- [Monitoring EBS R12.2 appsTier Provisioning Progress](#)
- [Preparing Target EBS R12.2 Environments for Provisioning](#)

- [Customizing Oracle VDB Environment Variables](#)
- [Provisioning an Oracle VDB](#)

Monitoring EBS R12.2 appsTier Provisioning Progress

This topic describes how to monitor the progress of EBS R12.2 appsTier provisioning.

The Delphix Engine automates the configuration of the appsTier during provisioning. It spends the majority of provisioning time running RapidClone utilities such as `adcfgclone`. The Delphix Engine does not report the progress of RapidClone utilities through the progress of the **DB_Provision** job. To track provisioning progress more accurately, Delphix recommends monitoring EBS log files generated on the target appsTier environment.

Procedure

1. Connect to the target appsTier environment using SSH or an alternative utility.
2. Change directories to the `<INST_TOP>/admin/log/`.
 - Replace `<INST_TOP>` with the value of `INST_TOP` on the virtual EBS instance.
3. After `adcfgclone` has begun running, a file matching `ApplyAppsTier_*.log` will exist. Identify this log file and use `tail` or an equivalent utility to monitor it.

Leave the Mount Point Prior to Refresh

If provisioning fails, you will need to fix the cause of the failure and perform a refresh of the dataset before you can attempt configuration again. Prior to refreshing or disabling the dataset, be sure to change directories to a location outside of the mount path on the target environment. If you leave a shell session with a current working directory inside the mount path, the dataset will fail to unmount cleanly during a refresh or disable.

Related Topics

- [Provisioning a Virtual EBS R12.2 Instance](#)

Monitoring EBS R12.2 dbTechStack Provisioning Progress

This topic describes how to monitor the progress of EBS R12.2 dbTechStack provisioning.

The Delphix Engine automates the configuration of the dbTechStack during provisioning. It spends the majority of provisioning time running RapidClone utilities such as `adcfgclone`. The Delphix Engine does not report the progress of RapidClone utilities through the progress of the **DB_Provision** job. To track provisioning progress more accurately, Delphix recommends monitoring EBS log files generated on the target dbTier environment.

Procedure

1. Connect to the target dbTier environment using SSH or an alternative utility.
2. Change directories to the `<ORACLE_HOME>/appsutil/log/<CONTEXT_NAME>/`.
 - Replace `<ORACLE_HOME>` with the path to the dbTechStack's Oracle Home. This path will be under the mount path specified during provisioning.
 - Replace `<CONTEXT_NAME>` with the virtual EBS instance's context name.
3. After `adcfgclone` has begun running, a file matching `ApplyDBTechStack_*.log` will exist. Identify this log file and use `tail` or an equivalent utility to monitor it.

Leave the Mount Point Prior to Refresh

If provisioning fails, you will need to fix the cause of the failure and perform a refresh of the dataset before you can attempt configuration again. Prior to refreshing or disabling the dataset, be sure to change directories to a location outside of the mount path on the target environment. If you leave a shell session with a current working directory inside the mount path, the dataset will fail to unmount cleanly during a refresh or disable.

Related Topics

- [Provisioning a Virtual EBS R12.2 Instance](#)

Managing Data Operations for Virtual EBS Instances

This section covers the following topics:

- [Starting and Stopping a Virtual EBS Instance](#)
- [Rewinding a Virtual EBS Instance](#)
- [Refreshing a Virtual EBS Instance](#)
- [Enabling and Disabling a Virtual EBS Instance](#)
- [Deleting a Virtual EBS Instance](#)
- [Modifying the appsTier Topology](#)
- [Replicating a Virtual EBS Instance](#)
- [Upgrading a Delphix Engine hosting a Virtual EBS Instance](#)
- [Migrating an EBS VDB](#)

Starting and Stopping a Virtual EBS Instance

This topic describes the process of starting and stopping a virtual Oracle E-Business Suite (EBS) instance.

You can start and stop virtual EBS instances through the **Delphix Management** application or through the standard Oracle command-line interface (CLI) utilities, `adstrtal` and `adstpall`. The Delphix Engine will show the dbTechStack and appsTier as running as long as there are processes using the dbTechStack and appsTier filesystem mounts on the target environments.

 Be careful! Services running on the appsTier depend on the availability of services on the dbTier. The steps below are explicitly ordered with these dependencies in mind. Executing steps out of order may lead to errors in accessing the virtual EBS instance.

Stopping

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **appsTier vFiles** for your EBS instance.

5. Click the **Stop** icon. 

Stopping the appsTier will run Oracle's `adstpall.sh` utility.

For multi-node appsTier, secondary nodes will be stopped sequentially, followed by the primary node.

 Stopping the appsTier may take a long time. The Delphix Engine will wait for all Oracle application processes to exit before declaring the appsTier as stopped.

6. Select the **VDB** utilized by your EBS instance.
7. Click the **Stop** icon.
This action will shutdown the database instance.
8. Select the **dbTechStack vFiles** hosting your virtual EBS database.
9. Click the **Stop** icon.
Stopping the dbTechStack will shutdown the database listener.

Starting

1. Login to the **Delphix Management** application using **Delphix Admin** credentials.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **dbTechStack vFiles** hosting your virtual EBS database.

5. Click the **Start** icon. 

6. Starting the dbTechStack will start the database listener.

7. Select the **VDB** utilized by your EBS instance.

8. Click the **Start** icon.

Starting the database will open the database.

9. Select the **appsTier vFiles** for your EBS instance.

10. Click the **Start** icon.

Starting the appsTier will run Oracle's `adstrtal.sh` utility.

For multi-node appsTier, the primary node will be started first, followed by secondary nodes sequentially.

Related Topics

- [Managing Data Operations for Virtual EBS Instances](#)

Rewinding a Virtual EBS Instance

This topic describes the process of rewinding a virtual Oracle E-Business Suite (EBS) instance.

 Changes applied to EBS and picked up only in some dSource snapshots may make certain combinations of snapshots across the appsTier and dbTier incompatible. When provisioning, refreshing or rewinding a virtual EBS instance, be sure the points in time you select for each dataset are compatible with each other.

 Be careful! Services running on the appsTier depend on the availability of services on the dbTier. The steps below are explicitly ordered with these dependencies in mind. Executing steps out of order may lead to errors in accessing the virtual EBS instance.

Prerequisites

The appsTier **Instance Home Directory** of the virtual EBS instance must reside under the specified **Mount Path**.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **appsTier vFiles** for your EBS instance.
5. Click the **Stop** icon to shutdown the appsTier services.
6. Select the **VDB** utilized by your EBS instance.
7. Click the **Stop** icon to shutdown the database.
8. Select the **dbTechStack vFiles** hosting your virtual EBS database.
9. Click the **Stop** icon to shutdown the database listener.
10. Rewind the dbTechStack vFiles.
 - a. Select a **snapshot**.
 - b. Click the **Rewind** button.
11. Rewind the EBS VDB.
 - a. Select a **snapshot**.
 - b. Click the **Rewind** button.
12. Rewind the appsTier vFiles.
 - a. Select a **snapshot**.
 - b. Click the **Rewind** button.

Once you have rewound all three EBS virtual datasets successfully, your virtual EBS instance should be running and accessible.

Related Topics

- [Managing Data Operations for Virtual EBS Instances](#)

Refreshing a Virtual EBS Instance

This topic describes the process of refreshing a virtual Oracle E-Business Suite (EBS) instance.

 Changes applied to EBS and picked up only in some dSource snapshots may make certain combinations of snapshots across the appsTier and dbTier incompatible. When provisioning, refreshing, or rewinding a virtual EBS instance, be sure the points in time you choose for each dataset are compatible with each other.

 Be careful! Services running on the appsTier depend on the availability of services on the dbTier. The steps below are explicitly ordered with these dependencies in mind. Executing steps out of order may lead to errors in accessing the virtual EBS instance.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**
4. Select the **appsTier vFiles** for your EBS instance.
5. Click the **Stop** icon to shut down the appsTier services.
6. Select the **VDB** utilized by your EBS instance.
7. Click the **Stop** icon to shut down the database.
8. Select the **dbTechStack vFiles** hosting your virtual EBS database.
9. Click the **Stop** icon to shut down the database listener.

Clean Up Might Be Required

If you did NOT specify the **Cleanup Before Provision** option for either your virtual dbTechStack or appsTier, you must manually clean up your target environments prior to refreshing. If you have specified this option for both datasets, no manual work is required.

To manually clean up a target environment prior to refresh, remove instance-specific directories and or inventory entries that will conflict with files and entries recreated during the refresh.

Without this clean-up, post-clone configuration performed during a refresh will fail with an error claiming that a conflicting EBS instance is already installed.

10. Refresh the dbTechStack vFiles.
 - a. On the vFiles, click the **Refresh**.
 - b. Select a **snapshot** from which to refresh.
11. Refresh the EBS VDB.
 - a. On the VDB, click the **Refresh**.
 - b. Select a **snapshot** from which to refresh.
12. Refresh the appsTier vFiles.
 - a. On the vFiles, click the **Refresh**.
 - b. Select a **snapshot** from which to refresh.

Once you have refreshed all three EBS virtual datasets successfully, your virtual EBS instance should be running and accessible.

Related Topics

- [Managing Data Operations for Virtual EBS Instances](#)

Enabling and Disabling a Virtual EBS Instance

This topic describes the process of enabling and disabling a virtual Oracle E-Business Suite (EBS) instance.

An enabled virtual EBS instance will be running and fully available to end-users. A disabled virtual EBS instance will be neither running nor mounted to the target environments.

 Be careful! Services running on the appsTier depend on the availability of services on the dbTier. The steps below are explicitly ordered with these dependencies in mind. Executing steps out of order may lead to errors in accessing the virtual EBS instance.

Disabling

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **appsTier vFiles** for your EBS instance.
5. From the Actions menu (...) select **Disable**.

Disabling the appsTier vFiles will stop the appsTier services and unmount the appsTier files.

 Stopping the appsTier may take a long time. The Delphix Engine will wait for all Oracle application processes to exit before declaring the appsTier as stopped.

6. Select the **VDB** utilized by your EBS instance.
7. From the Actions menu (...) select **Disable**.
Disabling the VDB will stop the database instance and unmount the data files.
8. Select the **dbTechStack vFiles** hosting your virtual EBS database.
9. From the Actions menu (...) select **Disable**.
Disabling the dbTechStack vFiles will stop the database listener and unmount the dbTechStack files.

Once you have disabled all three EBS virtual datasets successfully, your virtual EBS instance should be fully removed from the target environment.

Enabling

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **dbTechStack vFiles** hosting your virtual EBS database.
5. From the Actions menu (...) select **Enable**.
Enabling the dbTechStack vFiles will mount the dbTechStack files and start the database listener.
6. Select the **VDB** utilized by your EBS instance.
7. From the Actions menu (...) select **Enable**.
Enabling the VDB will mount the data files and start the database instance.
8. Select the **appsTier vFiles** hosting your virtual EBS database.
9. From the Actions menu (...) select **Enable**.
Enabling the dbTechStack vFiles will mount the appsTier files and start the application services.

Once you have enabled all three EBS virtual datasets successfully, your virtual EBS instance should be running and accessible.

Related Topics

- [Managing Data Operations for Virtual EBS Instances](#)

Deleting a Virtual EBS Instance

This topic describes the process of deleting a virtual Oracle E-Business Suite (EBS) instance.

⚠ Be careful! Services running on the appsTier depend on the availability of services on the dbTier. The steps below are explicitly ordered with these dependencies in mind. Executing steps out of order may lead to errors in accessing the virtual EBS instance.

Procedure

1. Login to the **Delphix Management** application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **appsTier vFiles** for your EBS instance.
5. Delete the appsTier vFiles by clicking the **Trash Can** icon in the lower left-hand corner.
6. Select the **VDB** utilized by your EBS instance.
7. Delete the VDB by clicking the **Trash Can** icon in the lower left-hand corner.
8. Click **Manage**.
9. Select **Environments**.
10. Select the **target dbTier environment**.
11. Click the **Databases** tab.
12. In the list of **Installation Homes** on the environment, click the **Trash Can** icon next to the dbTechStack Oracle Home you want to delete.
13. Click **Manage**.
14. Select **Datasets**.
15. Select the **dbTechStack vFiles** for your EBS instance.
16. Delete the dbTechStack vFiles by clicking the **Trash Can** icon in the lower left-hand corner.
17. Clean up any files that the virtual EBS instance might have created outside of the Delphix mount points on the target environments. These typically include the instance-specific directories, oraInventory files, and oraTab entries.

Once you have deleted all three EBS virtual datasets successfully, your virtual EBS instance should be fully removed from the target environments.

Related Topics

[Managing Data Operations for Virtual EBS Instances](#)

Modifying the appsTier Topology

This topic describes the process of modifying the appsTier topology of a virtual EBS instance.

appsTier Topology Changes Invalidate Rewind

If you modify the appsTier topology of a virtual EBS instance, all snapshots in existence prior to the modification will no longer be valid for rewind. The appsTier topology in the snapshots will no longer match the Delphix Engine's configuration and rewind targeting these snapshots will fail if attempted.

Procedure

1. Disable the virtual EBS instance by following the procedure outlined in [Enabling and Disabling a Virtual EBS Instance](#).
2. Select the **appsTier vFiles** for your EBS instance.
3. On the back of the card, modify the **EBS-specific parameters** for the virtual appsTier. This process will normally entail adding or deleting **Additional Nodes** and configuring its corresponding **Services**.
4. Apply the configuration changes by refreshing the entire virtual EBS instance. Follow the procedure outlined in [Refreshing a Virtual EBS Instance](#).

Related Topics

- [Managing Data Operations for Virtual EBS Instances](#)
- [Enabling and Disabling a Virtual EBS Instance](#)
- [Refreshing a Virtual EBS Instance](#)

Replicating a Virtual EBS Instance

This topic describes the process of replicating a virtual Oracle E-Business Suite (EBS) instance.

Replication

1. Login to the *source* Delphix Engine's **Delphix Management** application.
2. Configure replication between the source Delphix Engine and a target Delphix Engine.
For a detailed outline of the replication process, see [Configuring Replication](#).
3. Select the **dbTechStack vFiles**, **VDB**, and **appsTier vFiles** objects to be replicated.
These objects have dependencies on all other Delphix Engine objects relevant to the virtual EBS instance: you do not need to specify any additional objects for EBS replication.
4. Schedule or perform the replication.

Failover

1. Login to the *target* Delphix Engine's **Delphix Management** application.
2. Failover the replica **dbTechStack vFiles**, **VDB**, and **appsTier vFiles**. Failing over these object will sever future replication, but will not enable the datasets.
For a detailed outline of the failover process, see [Controlled Failover](#).
3. Enable the dbTier and appsTier environments.
 1. Click **Manage**.
 2. Select **Environments**.
 3. For each environment, from the Actions menu (...) select **Disable**.
 - Enable the virtual EBS instance by following the procedure outlined in [Enabling and Disabling a Virtual EBS Instance](#).

Provisioning from a Replicated EBS Instance

You can provision new virtual EBS instances from replicated datasets regardless of whether these datasets have been failed over. The provisioning process for each version of EBS is outlined in [Virtualizing Oracle E-Business Suite](#). For more information about provisioning from replicated datasets, see [Provisioning from Replicated Data Sources or VDBs](#).

Related Topics

- [Configuring Replication](#)
- [Provisioning from Replicated Data Sources or VDBs](#)
- [Enabling and Disabling a Virtual EBS Instance](#)
- [Virtualizing Oracle E-Business Suite](#)
- [Controlled Failover](#).

Upgrading a Delphix Engine hosting a Virtual EBS Instance

This topic describes the process of upgrading a Delphix Engine hosting a virtual Oracle E-Business Suite (EBS) instance.

Procedure

1. Disable the virtual EBS instance by following the procedure outlined in [Enabling and Disabling a Virtual EBS Instance](#).
2. Once you have safely disabled the virtual EBS instance, upgrade the Delphix Engine by following the procedure outlined in [Upgrading the Delphix Engine](#).
3. Enable the virtual EBS instance by following the procedure outlined in [Enabling and Disabling a Virtual EBS Instance](#).

Related Topics

- [Enabling and Disabling a Virtual EBS Instance](#)
- [Upgrading the Delphix Engine](#)

Migrating an EBS VDB

This topic describes how to migrate an EBS Virtual Database (VDB) from one target environment to another.

There may be situations in which you want to migrate a virtual database to a new target environment, for example when upgrading the host on which the VDB resides, or as part of a general data center migration. This is easily accomplished by first disabling the database, then using the Migrate VDB feature to select a new target environment.

There are three components for any EBS environment - dbTechStack, Database, and appsTier. Whenever we create virtual copies for any EBS environment, we provision these components in the following sequence: (1) dbTechStack (2) Database (3) appsTier.

We can **only** migrate the EBS database (VDB) with the following prerequisite.

Prerequisite

You should have already set up a new target environment that is compatible with the VDB that you want to migrate.

Procedure

1. Login to your Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **VDB** you want to migrate.
5. From the Actions menu (...) select **Disable**.
6. Click **Disable** to confirm.
When the VDB is disabled, its icon will turn grey.
7. From the Actions menu (...) select **Migrate**.
8. Select the new target environment for the VDB, the user for that environment, and the database installation where the VDB will reside.
9. Select **Migrate** to confirm your selections.
10. From the Actions menu (...) select the **Enable**
11. Click **Enable** to confirm.

Within a few minutes, your VDB will restart in the new environment, and you can continue using it.



- You need to [provision the EBS dbTechStack](#) on the new target environment with a similar DB name VDB that you want to migrate.
- EBS appsTier vFile cannot be migrated, you need to [re-configure the EBS appsTier](#) by modifying the custom configuration for Target DB server node information to the target environment where VDB is migrated.

Virtual EBS Instance Recipes

This section covers the following topics:

- [Configuring a Delphix Self-Service Data Template for EBS](#)
- [Managing the APPS Password](#)
- [Refreshing a Target dbTier Environment](#)
- [Managing the Weblogic Password](#)
- [Working with EBS Snapshots](#)

Configuring a Delphix Self-Service Data Template for EBS

This topic describes the process of configuring a Delphix Self-Service (Jet Stream) data template for Oracle E-Business Suite (EBS).

By configuring a data template for EBS, you eliminate the need to order operations applied across the EBS datasets. The ordering of operations is a result of the dataset dependencies discussed in topics under [Managing Data Operations of Virtual EBS Instances](#), such as [Refreshing a Virtual EBS Instance](#).

Procedure

1. Create a Delphix Self-Service data template by following the procedure outlined in [Understanding Data Templates](#).
 - a. For EBS, the data template will have three data sources: the dbTechStack, database, and appsTier.
2. Be sure to set the following **ordering of the data sources** when creating the data template. This ordering will ensure that Delphix Self-Service operations do not violate the EBS dataset dependencies.

Order	Dataset
1	dbTechStack
2	Database
3	appsTier

Once you have created a Delphix Self-Service data template, you can configure Delphix Self-Service data containers to manage virtual EBS instances. Delphix Self-Service data containers will follow the ordering of data sources configured in the template. All Delphix Self-Service operations should work as expected for virtual EBS instances.

Related Topics

- [Managing Data Operations of Virtual EBS Instances](#)
- [Refreshing a Virtual EBS Instance](#)
- [Understanding Data Templates](#)

Managing the APPS Password

This topic outlines how to manage the APPS password on a virtual EBS instance.

The Delphix Engine requires a virtual EBS instance's APPS password in order to manage its dbTechStack and appsTier. This password is encrypted when stored within the Delphix Engine and is exposed as an environment variable to RapidClone tools.

When a virtual EBS instance is being provisioned, the Delphix Engine's copy of the APPS password should match the source EBS instance's APPS password during its time of the snapshot. After provisioning, you can change the APPS password of the virtual instance in both EBS and the Delphix Engine.

Changing the APPS Password

1. Provision a virtual EBS instance.
2. Change the APPS password in the virtual EBS instance.
You can perform the password change using Oracle's `FNDCPASS` utility.
3. Stop the virtual appsTier manually using `adstpall`.
You cannot stop the virtual appsTier through the Delphix Engine, because the APPS password being stored is now out of sync.
4. Disable the entire virtual EBS instance.
For an outline of this process, see [Enabling and Disabling a Virtual EBS Instance](#).
5. Change the APPS password on the dbTechStack.
 - a. Select the **dbTechStack vFiles** hosting your virtual EBS database.
 - b. On the back of the card, click the **Custom** tab.
 - c. Edit the **APPS Password** field.
6. Change the APPS password in any hook operations defined on the EBS virtual database.
 - a. Select the **VDB** utilized by your EBS instance.
 - b. On the back of the card, click the **Hooks** tab.
 - c. Edit the relevant hook operations.
Typically, you will need to edit the Pre-Snapshot hook operation running `adpreclone`.
7. Change the APPS password on the appsTier.
 - a. Select the **appsTier vFiles** hosting your virtual EBS database.
 - b. On the back of the card, click the **Custom** tab.
 - c. Edit the **APPS Password** field.
8. Enable the entire virtual EBS instance.
For an outline of this process, see [Enabling and Disabling a Virtual EBS Instance](#).

Refreshing or Rewinding

The APPS password stored across individual snapshots of a virtual EBS instance will not be consistent after a password change. Old snapshots of EBS data will refer to a different APPS password than new snapshots of EBS data. To perform a refresh or rewind, you must explicitly manipulate the Delphix Engine's copy of the APPS password to ensure that the virtual EBS instance is being accessed with the correct APPS password at every step.

Recipe Not Needed If Password Not Changed

These steps are only necessary if the virtual EBS instance has a different APPS password than the snapshots being targeted by the refresh or rewind.

If the APPS password has not been changed, follow the instructions in [Refreshing a Virtual EBS Instance](#) or [Rewinding a Virtual EBS Instance](#).

1. Before refreshing or rewinding the virtual EBS instance, disable the entire virtual EBS instance.
For an outline of this process, see [Enabling and Disabling a Virtual EBS Instance](#).
2. Identify the APPS password for the snapshots being targeted by the refresh or rewind. Modify the virtual EBS instance to refer to this password.
 - a. Change the APPS password on the dbTechStack.
 - i. Select the **dbTechStack vFiles** hosting your virtual EBS database.
 - ii. On the back of the card, click the **Custom** tab.
 - iii. Edit the **APPS Password** field.
 - b. Change the APPS password in any hook operations defined on the EBS virtual database.
 - i. Select the **VDB** utilized by your EBS instance.
 - ii. Select the Configuration tab, then select the **Hooks** tab.
 - iii. Edit the relevant hook operations.

Typically, you will need to edit the Pre-Snapshot hook operation running `adpreclone` and the Configure Clone operation running `adcfgclone`.
 - c. Change the APPS password on the appsTier.
 - i. Select the **appsTier vFiles** hosting your virtual EBS database.
 - ii. Click the **Custom** tab.
 - iii. Edit the **APPS Password** field.
3. Perform the refresh or rewind while the EBS instance is disabled.
See [Refreshing a Virtual EBS Instance](#) or [Rewinding a Virtual EBS Instance](#) for an outline of these processes.

Related Topics

- [Enabling and Disabling a Virtual EBS Instance](#)
- [Refreshing a Virtual EBS Instance](#)
- [Rewinding a Virtual EBS Instance](#)

Refreshing a Target dbTier Environment

This topic describes how to properly refresh a target dbTier environment hosting an EBS VDB.

When an environment is initially registered with the Delphix Engine, a discovery process is responsible for registering resources such as Oracle Homes, listeners, and databases. When an environment is modified – for example, when an Oracle Home is upgraded or removed – you must run the discovery process again in order to register new resources and unregister removed ones.

The virtual dbTechStack plays a unique role in the Delphix Engine because it acts both as a virtual dataset AND a discovered resource on the target dbTier environment. Because of this unique role, you must mount and start this dataset (hosting the EBS database listener) before you refresh the target dbTier environment.

If you unmount or stop the virtual dbTechStack, the discovery process will assume that the EBS Oracle Home and database listener have been removed from the environment. Further interactions with the EBS VDB will result in an error, because neither the database listener nor Oracle Home appears available to the database.

Procedure

1. Prior to refreshing a target dbTier environment, ensure that the virtual dbTechStack is both enabled and started.
For more information, see [Enabling and Disabling a Virtual EBS Instance](#) and [Starting and Stopping a Virtual EBS Instance](#).
2. Refresh the target dbTier environment.

Related Topics

- [Enabling and Disabling a Virtual EBS Instance](#)
- [Starting and Stopping a Virtual EBS Instance](#)

Managing the Weblogic Password

Changing the Weblogic Password

For existing EBS 12.2 customers who intend to use this WebLogic password change feature need to follow the below steps for their EBS application:

1. Disable all the virtual datasets created on Delphix Engine in the following order:
 - a. appsTier vFiles
 - b. VDB utilized by your EBS instance
 - c. dbTechStack vFiles
2. Disable all the EBS dSource created on Delphix Engine once you have disabled all three EBS virtual datasets in the above step.
3. Upgrade Delphix Engine to 5.3.2.0 release and upload EBS 1.4.0 plugin release.
Go to Manage>Plugins
 - Upload a plugin file to upgrade the existing plugin by clicking plus (+) symbol.
4. Refresh all host environments added on the Delphix Engine.
Go to Manage> Environments
 - From the symbol (...) located in the upper left-hand corner, select Refresh All.
5. Change the Weblogic admin password on EBS appsTier vFile.
 - Select the appsTier vFiles hosting your virtual EBS database.
 Go to Configurations >Custom
 - Update the **Target Weblogic AdminServer Password** field with the password you would like to set for EBS appsTier.
6. Enable all the EBS dSource created on Delphix Engine.
7. Enable all the virtual datasets created on Delphix Engine in the following order:
 - a. appsTier vFiles
 - b. VDB utilized by your EBS instance
 - c. dbTechStack vFiles
8. Verify the changed WebLogic password by logging into the WebLogic console.

EBS 12.2 customers configuring appsTier on a multi-node environment will need to manually enable passwordless SSH on all target nodes *only when* they want to change the Weblogic password and then proceed with EBS appsTier provision. The passwordless SSH **should** be enabled between all nodes and for both low privileged and high privileged user.

Enabling Passwordless SSH between target Application Nodes

To enable passwordless SSH between target application nodes complete the following

The following is an example of a two-node EBS appsTier.

Node 1: test1-ap1

Node 2: test1-ap2

1. Create Authentication SSH-Keygen Keys on Node 1.

```
[applvis@test1-ap1 ~]$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/applvis/.ssh/id_rsa):
Created directory '/home/applvis/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/applvis/.ssh/id_rsa.
Your public key has been saved in /home/applvis/.ssh/id_rsa.pub.
The key fingerprint is:
b8:1d:83:ef:40:d3:54:d6:50:2d:74:28:66:8a:4d:15 applvis@test1-ap1.dcenter.delphix.com
The key's randomart image is:
+---[ RSA 2048]-----+
|           .E*oo. |
|           .o+ +.. |
|          +.+ . . |
|           .=o    |
|           = S    |
|           . = o   |
|           o o    |
|           o      |
|           .      |
+-----+
[applvis@test1-ap1 ~]$
```

2. Create .ssh Directory on Node 2.

```
[applvis@test1-ap1 ~]$ ssh applvis@test1-ap2 mkdir -p .ssh
The authenticity of host 'test1-ap2 (10.43.17.170)' can't be established.
RSA key fingerprint is 22:83:2b:d2:53:03:3e:b6:8b:fe:6a:87:21:35:d5:fc.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'test1-ap2,10.43.17.170' (RSA) to the list of known hosts.
applvis@test1-ap2's password:
[applvis@test1-ap1 ~]$
```

3. Upload Generated Public Keys to Node 2.

```
[applvis@test1-ap1 ~]$ cat .ssh/id_rsa.pub | ssh applvis@test1-ap2 'cat >> .ssh/authorized_keys'
applvis@test1-ap2's password:
[applvis@test1-ap1 ~]$
```

4. Set Permissions on Node 2 and Login to verify without password.

```
[applvis@test1-ap1 ~]$ ssh applvis@test1-ap2 "chmod 700 .ssh; chmod 640 .ssh/authorized_keys"
[applvis@test1-ap1 ~]$
[applvis@test1-ap1 ~]$ ssh test1-ap2
Last login: Mon Jan 28 21:46:12 2019 from 10.43.15.129
[applvis@test1-ap2 ~]$
[applvis@test1-ap2 ~]$ hostname -s
test1-ap2
```

Related Topics

- [Virtual EBS Instance Recipes](#)

Working with EBS Snapshots

Taking a snapshot creates a new snapshot entry in the dSource's Timeflow. The plugin-managed dSource allows you to select the parameters before taking the snapshot. These parameters are provided as an input to pre-snapshot and post snapshot functions. On the other hand, in the default snapshot, the parameters are pre-defined in the dSource environment.

Snapshot (Default)

This section lists the steps to take a snapshot and delete the same.

1. Login to the **Delphix Management** application.
2. Click **Manage** and select **Datasets** from the dropdown list.
3. Select the dSource you want to Snapshot.
4. Click the **Camera** icon. Alternatively, click the arrow next to the Camera icon and select **Snapshot (default)**.
5. From the Snapshot dialog, select **Perform Snapshot**.
6. Click **View:** under **Timeflow** to verify the Snapshot you just created.
7. To delete, select the snapshot you just created, and from the Actions menu (...) select **Delete Snapshot**.
8. In the Delete Snapshot dialog, select **Delete**.
9. Click **View:** under **Timeflow** to verify the Snapshot, you just deleted.

Best Practices

For optimal usage of Delphix, there are many different configurations and best practices you'll need to note to tune your engines for peak performance. This ranges from settings on the hypervisor that runs Delphix to the setting on your sources, virtual databases, and the operating systems on which they run. Find out more information based on each category below.

These pages are the starting point to begin your architecture planning for the best possible Delphix deployment. To function as a high-performance virtual appliance, we need to ensure the underlying infrastructure components are deployed in an ideal manner consistent with our best practices - these pages are to help you achieve that goal.

Introduction to Delphix Architecture

[Database Virtualization with Delphix](#)

Hypervisor and Host

[Best Practices for Hypervisor Host and VM Guest](#)

Network

[Best Practices for Network Configuration](#)

[CLI Cookbook: Network Performance](#)

[Optimal Network Architecture for the Delphix Engine](#)

[Network Operations Using the Delphix Session Protocol](#)

[Network Performance Expectations and Troubleshooting](#)

Storage

[Best Practices for Storage](#)

[Optimal Storage Configuration Parameters for the Delphix Engine](#)

[Storage Performance Expectations and Troubleshooting](#)

Data Protection

[Best Practices for Delphix Engine Data Protection](#)

Source DB and OS Settings

[Best Practices for Source DB and OS Settings](#)

Target DB and OS Settings

[Best Practices for Target DB and OS Settings](#)

Staging Target

[Best Practice for Staging Targets](#)

Architecture Checklist (Consolidated Best Practices)

[Architecture Checklist](#)

Architecture Checklist FAQ

[Architecture Checklist FAQ](#)

Architecture Checklist FAQ

Architecture Best Practices Overview FAQ

These questions relate to the standard processes our Solutions Architecture team executes when engaged for Architecture review of a customer's environment.

Why go through an architecture or sizing process?

To measure - when possible - the workload in your environment and calculate the number of engines needed

- To ensure best practices are applied so your engine(s) perform to the level expected
- To find and avoid or work around potential challenges in an environment
- To maximize your license value

Why does Delphix collect IO data?

The collection of IO data allows us to more accurately size the Delphix Engine rather than making recommendations based on assumptions or synthetic data. The data you provide for analyses is from the SOURCE and TARGET DB servers.

Why is the Customer Database and Inventory (CDI) important?

The database information is typically the fundamental building block upon which we build our understanding of your future Delphix environment. We ask questions related to the database names, locations, platforms, and versions, as well as estimates of size, throughput, and changes.

The answers you provide to our infrastructure questions give us insight into potential challenges that may arise as we seek to integrate our application and best practices into your environment. Processes and software you use normally may not integrate optimally into our engine, and it's helpful to know this as soon as possible.

Why does Delphix use a Virtual Machine? Shouldn't all high-performance applications use physical infrastructure?

Using virtualization not only allows our customers to use their choice of commodity hardware, but it also allows us to focus on core features rather than hardware support.

Architecture Best Practices for Hypervisor Host and VM Guest - ESX FAQ

Why does Delphix recommend 8 vCPUs and 128GB of memory per 8 vCPUs?

8 vCPUs are not only a standard licensing block, but they are also key to meeting our 10Gbps single-engine throughput potential and help to sustain low latency for VDBs.

As with CPU, cache memory is required to drive peak loads on the Delphix Engine. More memory allows for more blocks to be read from the cache rather than going to less performant disks. Delphix stores cached data in a compressed format and only keeps a single copy of unique blocks in memory. These features give read performance across multiple VDBs provisioned from a single source dramatic improvements in speed, scalability, and memory utilization.

Why does Delphix request reservations for CPU and memory?

Delphix performance can be greatly impacted there is contention over CPU or RAM. Reservations allow the engine to explicitly control those resources and avoids the possibility of contention with other VMs, even when resources are overcommitted.

Why does Delphix request Hyper-Threading (HT) be disabled?

Hyper-Threading can have a positive impact on many, but not all applications. Delphix has a different execution profile that does not benefit from Hyper-Threading. Factors such as significant memory bandwidth requirements and a high level of parallelism which requires a high number of shared locks mean that HT and Delphix do not generally work well together.

Why does Delphix request 4 controllers, and why must the storage be identical between them?

To provide optimal storage performance, you must spread data equally over the maximum (4) virtual SCSI controllers. To provide consistent performance between each of the four controllers, you need to ensure storage is identical between them.

Why must virtual disks (VMDKs) be thick provisioned and eager-zeroed?

Thick provisioning and eager zero ensure performance is top-notch from the start with no hiccups from expanding virtual resources.

Why is 20% free space required?

While the ZFS file system has a lot of features leveraged by Delphix, it loses efficiency as space decreases. 20% is the minimum that must be available for best performance.

Why does Delphix request you reserve CPU and RAM for Hypervisor overhead?

Based on VMWare's resource management guide and our own experiences with high IO throughput. Note there is no specific mechanism to assign resources to the hypervisor, the only way to preserve overhead is by not allocating resources to guests.

Why does Delphix generally want VMWare HA enabled, but DRS disabled?

VMware HA (High Availability) addresses outages that occur when a physical host goes down or is completely offline, by migrating the guest(s) to another physical host and restarting them. There is no real downside, it simply brings unavailable servers back online.

VMware DRS (Distributed Resource Scheduler) is for load balancing host resources in a cluster. Because of high IO and best practices configuration for optimal performance, our engine is typically not a good candidate for relocation.

Why does Delphix request you set power management to High-Performance Mode?

This will ensure power management will never impact performance by entering into a lower power state (also known as c-state).

Architecture Best Practices for Storage FAQ

Why does Delphix require 127GB of storage for the OS?

The system partition requires space to store and operate the OS, as well as application logs, upgrade and rollback images, and enough free space to store a kernel or application core dump should it be required.

Why does Delphix require our LUNS to be uniform and contain an equal quantity and capacity of VMDKs, yet thin provisioning is OK?

Because our engine leverages parallel reads, we need the storage capacity and quantity of disks they hold to be consistent. This allows the reads and writes to be evenly distributed, and minimizes the impact of potential utilization imbalances which would create a “long tail” of higher latency on a single controller, impacting the entire engine.

Data storage LUNS are generally formatted with a VMFS file system and have placed upon them a virtual disk (VMDK) which is thick provisioned and eager zeroed, so it would be a waste of time to thick provision the LUN also.

Why does Delphix require < 10ms latency (95th percentile) storage?

Storage latency is especially important in database environments. Average latency doesn't give a complete picture of responsiveness, especially because Delphix leverages parallel reads; so inconsistent performance (e.g. good average latency but a “long tail”) can impact multiple operations. This is why Delphix has a focus on 95th percentile latency, and why we validate storage performance as the first step when a new engine is deployed.

For more information, you may find the following article helpful: [The Cost of Latency](#).

Why does Delphix prefer to extend existing storage rather than simply add more while maintaining equal distribution?

While it is possible to add more storage and maintain the practice that "storage should be equal across controllers" – extending LUNS (then virtual disks) ensures that:

- We do not continue to fill disks which may be full
- Existing disks do not suffer a write performance penalty from low capacity
- Storage performance is consistent

Architecture Best Practices for Network FAQ

Why does Delphix request 10GE Ethernet?

As a matter of physics and standards - 10 gigabit (Gb) Ethernet can sustain approximately 1 gigabyte (GB) per second of throughput. With all our best practices applied, a Delphix Engine can achieve very close to that line speed, allowing for optimal load, engine, and license utilization. Lower network speeds may be acceptable for low loads, while in some environments NIC teaming (e.g. LACP) may be required for top speeds.

Why does Delphix require < 1ms latency to TARGET servers and < 50ms to SOURCE servers?

Delphix leverages NFS and iSCSI (depending on platform) for live TARGET DB mounting over the network, so it's imperative that latency is as low as possible. Data coming from SOURCE servers is not generally as time-sensitive, so you need a minimum latency of < 50ms to ensure operational integrity.

Why does Delphix request Jumbo Frames?

Jumbo frames increase the Ethernet maximum transmission unit (MTU) from the default 1500 bytes to 9000 bytes. This has several effects such as decreasing CPU cycles by transferring fewer packets and increasing the engine throughput. You will find jumbo frames have a 10-20% real-world impact and are required (along with all other best practices) to handle peak loads of 800 - 1000MB/s on an 8 vCPU engine with a 10Gb network.

How does Delphix avoid communication impact with non-jumbo frame hosts when Jumbo Frames are enabled on the Delphix Host?

[Path MTU Discovery](#) is the mechanism by which two hosts agree on the MTU leveraged for communication between them. This mechanism will ensure communication between both standard and Jumbo Frame enabled hosts works as expected.

When does Delphix recommend NIC teaming?

The Delphix Engine is capable of high throughput, but not every enterprise has sufficient network bandwidth to support it. Teaming is a less expensive way of increasing the bandwidth when compared to new hardware.

Why does Delphix recommend logical and physical and co-location?

The Delphix Engine leverages network connections extensively, so optimizing the latency whenever possible is very important - sometimes critical.

Architecture Best Practices for Delphix Engine Data Protection FAQ

Why does Delphix recommend SAN snapshots or Delphix replication for backup?

There are a few possible methods for data protection of the Delphix Engine. Those methods are SAN snapshots, Delphix replication, and virtual machine snapshots (for very small engines only). Because the Delphix Engine is itself a backup of source environments, many customers simply plan to rebuild in the event of a disaster.

What is the DXToolkit, and how can it help?

The professional services team has created a Perl-based “DXToolkit” which can help export and import certain configuration data over web services. This toolkit can be leveraged to assist with would normally be a manual re-install outside of the above methods.

For further detail around data protection, please speak with your Delphix contact.

Can I use a VMware-based backup solution such as VEEAM to backup my Delphix Engine?

Yes, VMWare backup solutions are useful for backing up guest VMs. However, Delphix suggests that you only use this approach for Delphix Engines which have a smaller storage footprint (perhaps < 2 TB) and are less active.

Running this type of backup puts a load on the environment, which might adversely impact Delphix VM performance.

Can I use a VMware snapshot for backing up Delphix for a small window – for example, during an engine upgrade?

Yes. However, even though snapshots are instantaneous, they track changes separately from the base disks and can grow to consume as much space as the original.

Upgrades, in particular, can change substantial amounts of data.

If you lose physical disks, snapshots are useless because it needs them to make up the current state of VM.

A Delphix Engine is often allocated multiple terabytes of storage and is often very busy due to load aggregation from virtual databases on multiple target servers, so this approach may be challenging.

Snapshots cannot detect storage corruption.

Can I use a Storage snapshot solution to protect Delphix against Storage and Delphix corruption?

Yes. However, please note that the caveats which apply to VMWare snapshots will also apply here.

A specific concern related to storage layer snapshots is that you must create a consistency group that contains both the OS and Data disks.

Can I use RMAN to backup my VDBs just like a Physical database to provide extra protection?

You can backup Delphix VDBs using Oracle RMAN tools, but the recovery database would first require re-hydration of that VDB, which might take up equivalent production storage space.

Furthermore, that re-hydrated database needs to be brought into the Delphix framework as a dSource, after which you can provision a VDB to complete recovery. The whole process might take hours or days to recover.

The best approach is to use the VDB Snapshot capability to backup VDB frequently and then leverage Delphix Replication capability to protect underlying Delphix storage, which holds that VDB snapshot.

Related Topics

- [Architecture Checklist](#)
- [Best Practices](#)

Architecture Checklist

Overview

The Architecture Checklist is an overview combining the more important Architectural Best Practices into a single list. For more detail regarding some of the reasons for our best practices, review the Architecture Checklist FAQ.

Architecture Best Practices for Hypervisor Host and VM Guest - ESX

Hypervisor

1. ESXi 6.x or 7.0 is recommended. ESXi 5.5 or earlier is no longer supported.
2. HyperThreading (HT) for Intel®-based servers (no HT on AMD CPUs).
 - a. Disable HT in BIOS, on the ESXi Host, **and** disable [HT Sharing](#) on the Delphix VM for consistency. This is our best practice, disable at all levels.
Any other combination may result in non-deterministic performance.
 - b. When HT cannot be turned off for both the Host and Delphix VM, it should be turned on at all levels, not run in a "mixed-mode".
 - i. HT disablement at a guest level only can result in non-deterministic performance.
 - ii. A dedicated ESXi host, cluster or DRA is recommended where consistent VDB performance is paramount.
 - c. VM migration to a new host (e.g. VMware HA or vMotion) can create mismatched HT settings.
3. ESXi overhead (resources required for hypervisor cannot be reserved, they must be left unallocated).
 - a. Memory Overhead - 10% of available RAM must not be allocated to guest VMs.
 - *Example: 256GB RAM, allocate 230GB to Delphix VM, leave 26GB for ESX*
 - b. CPU Overhead - At least 2 cores (ideally 4) must not be allocated to guest VMs.
 - *Example: If 16 physical cores are available, allocate 12 to the virtual machines, leaving 4 for the hypervisor*
 - Why 4? Certain hypervisor functions require precedence over any virtualized system. If a hypervisor needs more CPU than the amount currently available, it can de-schedule all other virtual processes to ensure adequate CPU resources for the hypervisor. Ensuring the hypervisor will not have to de-schedule any running virtual processes (worlds) by setting aside and not over-subscribing CPUs for virtual functions will leave them available for hypervisor use.
 - c. Even if the Delphix VM is the only VM on a host, the hypervisor is still active and essential; and still needs resources.
4. BIOS Power Management should be set to High Performance where ESXi controls power management.
 - a. Can be impacted by [VMware KB 1018206](#) - poor VM application performance caused by power management settings.
 - b. Ensure that all BIOS managed C-States other than C0 are disabled if power management is hardware controlled.
 - c. Ensure that all ACPI sleep states above S0 are disabled in the BIOS.
 - d. Examples for popular server lines from Cisco, HP, Dell below. *Specific models will vary, use the appropriate spec sheet.*
 - i. **UCS**: disable the Processor Power States, disable Power Technology, set Energy Performance to "Performance"
 - ii. **HP Proliant**: set HP Power Regulator to HP "Static High Performance" mode
 - iii. **Dell**: set BIOS System Profile to "Performance Optimized" mode
5. VMware HA can be enabled; VMware DRS is generally disabled.
6. Blade/Rack Server Firmware and ESXi Drivers should be updated to latest versions.

7. For Intel®-based servers with E5-2600 v2 processors.
8. Two typical server configurations:
 1. Blade Farm
 2. Rack Server
 Hyper-Converged configurations are possible for high performance.

Virtual Machine Guest

1. For VM machine settings, see [Virtual Machine Requirements for VMware Platform](#).
2. VMWare Guest Specifications:
 - a. Minimum: 8 vCPU x 64 GB
 Small: 8 vCPU x 128GB
 Medium: 16 vCPU x 256 GB
 Large: 24 vCPU x 512 GB
 - b. Reserve 100% of RAM and CPU:
 - If the ESX host is dedicated to Delphix, CPU and RAM reservations are advised but not necessary, however, swap space will be required on the hypervisor to compensate for the lack of reserved RAM.
 - c. [Hyperthreading](#) - See the ESX host section at the top. Disable HT Sharing on VM, disable HT on ESX Host.
3. Assign single-core sockets for vCPUs in all cases. If there is a compelling reason to use multi-core CPUs, reference the following article from VMware which describes matching virtual multi-core sockets to the hardware ESX is running on.
[VMware Article on CoresPerSocket](#)
 Example: *ESXi Host has 2 socket x 18 core Intel Xeon, Delphix Engine wants 16 vCPU.*
 Configure Delphix VM with 2 Virtual Sockets, 8 Cores Per Socket to utilize hardware architecture.
4. Avoid placing other extremely active VMs on the same ESX host.
5. Monitoring - vSphere Threshold Alerts for CPU, Network, Memory, Capacity.
6. To set the number of vCPUs per virtual machine via the vSphere client, please see "[Virtual CPU Configuration](#)" in the Administration guide:
[ESXi 6.0](#)
 - a. [Delphix VM CPU Utilization](#) - Delphix KB article on what makes Delphix VMs similar to other resource-intensive applications
 - b. [Exchange on VMware Best Practices](#) - VMworld 2013 session
 - c. [ESXTOP Reference, Blog](#)
7. Ensure that the latest available VMware drivers and firmware versions are installed for HBAs, NICs and any other hardware components configured on the Delphix virtual machine. This is a critical step that can have a massive impact on the performance and robustness of our solution.

Architecture Best Practices for Network Configuration

1. For more information about network configuration refer to [Network Performance Configuration Options](#).
2. Delphix Engine <====> Target Host:
Implement standard requirements for optimal NFS/ISCSI performance:
 - a. Optimal physical network topology:
 - Low latency: < 1ms for 8K packets.
 - Network adjacency: minimize network hops, co-locate in the same blade enclosure, co-locate on the same physical host.
 - Eliminate all Layer 3+ devices - firewalls, IDS, packet filters (Deep Packet Inspection - DPI).
 - Multiple switches can add latency and fragmentation and reordering issues will add significant latency.
 - b. Optimal throughput:

- 10GbE physical uplinks or higher.
 - Jumbo frames (typically MTU 9000) improves network efficiency for 8K packets: lower CPU, lower latency, greater throughput.
 - All devices end-to-end *must* be configured for the larger frame size including switches, routers, fabric interconnects, hypervisors and servers.
 - Traffic between two VMs on the same host is limited to ~16Gbps when leveraging built-in virtual networking.
- c. Optimal logical flow:
- Disable QoS throttles limiting network utilization below line rate (e.g. HP Virtual Connect FlexFabric).
 - Consider a dedicated VLAN (with jumbo frames) for NFS/iSCSI traffic.
 - [VMware KB-Configuring iSCSI port binding with multiple NICs in one vSwitch for VMware ESXi 6.0.x \(2045040\)](#)
 - [VMware KB-Considerations for using software iSCSI port binding in ESX/ESXi \(2038869\)](#)
- d. NIC Teaming (at ESX layer) of multiple physical uplinks can provide additional throughput for higher workloads.
- i. Examples: 4 x 1Gb NICs support up to 400 MBPS IO, 2 x 10Gb NICs support up to 2 GBPS IO.
 - ii. [VMware KB-1004088](#) has NIC teaming recommendations, including `route-based-on-IP-hash` policy.
 - iii. [VMware KB-1001938](#) has host requirements for physical link aggregation (LACP, EtherChannel).
 - iv. [VMware KB-1007371](#), [popular blog post](#) details problems with NIC selection using `dest-IP hash`.
- e. Fragmentation and dropped packets can result in excessive retransmissions of data, reducing throughput.
- i. On AIX, LSO (Large Send Offload) and LRO (Large Receive Offload) network features have caused many problems. The virtualization engine employs LSO only, but no LRO, AIX employs both. These features are enabled by default. The best practice is to disable LSO on the VE when the target is AIX.
 - LSO enabled on the Virtualization engine can result in dramatically limited or blocked transmit throughput (and increase re-transmission traffic).
 - When wishing to disable all, LSO should be disabled on the Delphix Virtualization Engine, while both LRO and LSO should be disabled on guest VM of AIX Target. The specific commands to disable LSO and LRO on AIX are not captured here and appear to be context/version-specific.
- f. **Jumbo frames check via ping**

```
Delphix Engine
$ ping -D -s [Target_IP] 8000
"ICMP Fragmentation needed and DF set from gateway" indicates MTU < 8028
```

```
Linux
$ ping -M do -s 8000 [Delphix_Engine_IP]
"Frag needed and DF set (mtu = xxxx)" indicates MTU < 8028
```

```
MacOS
ping -D -s 8000 [Delphix_Engine_IP]
```

Note: "sudo sysctl -w net.inet.raw.maxdgram=16384" will increase the max ICMP datagram size on Mac, allowing you to use -s 9000 on Mac0

```
Windows
ping -f -l 8000 [Delphix_Engine_IP]
```

<http://www.mylesgray.com/hardware/test-jumbo-frames-working/>

- g. Measure Network Bandwidth and Latency:
 - i. Latency in both directions should be < 1ms for an 8KB payload.
 - ii. Network Hops should be minimized: `tracert` (Unix/Linux) / `tracert (windows)`.
 - iii. Throughput in both directions: 50-100 MB/s on 1 GbE, 500-1000 MB/s on 10 GbE physical link.
 - iv. Always use the CLI when possible.
- h. NIC should use Auto-negotiate on Ethernet with a minimum of 1000Mbps.
 - Hard setting speed/duplex will limit network throughput below the line rate.
3. **Delphix <====> Staging Server** (SQL Server, Sybase):
Measure latency, bandwidth for transaction log restore performance
4. **Source <====> Delphix:**
Measure latency, bandwidth to explain snapsync performance
5. **ESX host <====> ESX host** (ESX Cluster):
 - i. This is because the entire memory footprint of the Delphix VM (more precisely, the entire address space of the ESX processes that comprise the VM) must be copied to the receiving ESX host, along with all changes to that address space as they happen. Our ZFS cache comprises the bulk of that memory and changes as I/O occurs, and that rate of change is governed by the network bandwidth available to the Delphix Engine.
 - ii. Attempting to live vMotion a DE with a 10Gb NIC over a 1Gb vMotion network is thus likely to fail.

Architecture Best Practices for Storage

1. Configuration:
 - a. Virtual Disks (VMDK) with spinning or tiered media must be thick provisioned + lazy zeroed.
 - For storage which is 100% SSD/EFD/Flash-based, continue to thick provision, however, eager zero is not necessary.
 - b. VMDKs may be homed on VMFS or RDM storage, however, VMFS is much more common and generally preferred.
 - c. Regardless of VMFS/RDM selection, physical LUNs must have uniform characteristics (RAID, spindle count, tier) and should be thin provisioned to save time.
 - d. Storage allocated must be identical between each vSCSI controller.

- e. The supported maximum of four virtual SCSI controllers (PVSCSI (default) or LSI Logic Parallel) must be enabled. A mix of different types of SCSI controllers is not supported within the engine.
 - f. Virtual Disks must be evenly distributed across the 4 virtual SCSI controllers. For example, 8 virtual disks should be configured as 2 disks per controller: SCSI (0:0), SCSI (0:1), SCSI (1:0), SCSI (1:1), SCSI (2:0), SCSI (2:1), SCSI (3:0), SCSI (3:1).
 - You don't need to account for the OS in the even distribution of disks across controllers, just pick one; the OS doesn't place a substantial load on the controller.
 - g. To provision VMDK disks over 16TB in size, the vSphere web client must be used, the win32 client will return an error.
 - h. As of 5.1.3, we require 127GB for the system disk. older versions defaulted to 150GB.
 - Due to our unified OVA for masking, the 127GB requirement also applies to Masking Engines
 - i. VMDK for the Delphix Engine OS is often stored on the same VMFS volume as the Delphix VM definition file (aka VMX). In that case, the VMFS volume must have sufficient space to hold the Delphix VMX Configuration, the VDMK for the system disk, a swap/paging area if the memory reservation was not enabled for the Delphix Engine (or to suspend the system), and any related VMware logging.
 - j. Set ESX Storage Multipathing IO optimization ([KB 2072070](#)).
 - i. Set multipathing policy to round-robin.
 - ii. Set path switching IO operation limit to 1 (default 1000).
 - k. Verify that the queue depth setting on the virtual disks used for VMware is appropriate based on the minimum of the HBA type or the underlying storage's combined queue depth for all hosts attached to the same controller
 - See this [VMware KB article](#) for how to check/set the queue depth
2. Testing:
- a. Run Storage Performance Tool on the raw storage **before any engine configuration**. This is a one-time opportunity for each engine upon installation
 - b. Required maximum storage latency is < 2ms for writes and < 10ms (95th percentile) for small random reads.
Minimum passing grades: 4KB/8KB reads (B-), 1MB reads (A), 1KB/128KB writes (A).
 - c. If working with the Delphix Professional Services team, we would expect to run additional baseline performance measurements via composite tools and scripts.
 - e.g. "Sanity Check" (Oracle) or "DiskSpd" (MSSQL)
3. Detail Discussion:
- a. Before beginning any discussion on storage performance, or at the beginning of the installation, collecting the following specs from your storage administrator will assist in understanding
 - b. Vendor, Model (For example: EMC VMAX 40k, HP 3PAR StoreServ 7000)
 - c. IO latency SLO (For example 5ms 99.999%)
 - d. IOPS/GB SLO (For example: 0.68 IOPS/GB for EMC Gold-1)
 - e. Cache type and size (For example FAST cache 768GB)
 - f. Tier, #Pools; if auto – tiering; relocation schedule (For example: Gold/Silver/Auto/3 pools/etc)
 - g. Pool detail: (#) drives, RPM, Type (For example: Pool1: (20) EFD, (30) 15k SAS, Pool 2: (40) 10k SATA)
 - h. Connection (For example: 16Gb Fibre Channel, 10Gb iSCSI, 20Gb N

Architecture Best Practices for Delphix Engine Data Protection

1. For protection against physical host failure - leverage VMware HA.
2. For protection against storage failure - leverage Delphix replication.
3. For protection against administrative error - leverage storage snapshots.
4. For protection against site failure - leverage Delphix replication and/or Delphix Live Archive.
5. Infrastructure Backup of the Delphix VM:
Must take a consistent group snapshot of all Delphix VM storage (system disk, VM configuration, database VMDK/RDMs)
RTO, RPO are inferior compared to Delphix Replication. The granularity of restore is at the VM level: **all-or-nothing**.

- a. Virtual Machine Backup: Create VM snapshot, backup (proxy server), remove the snapshot.
Products use VMware APIs: NBU for VMware, TSM, NetWorker
Limited to < 2TB because of time to backup, impact on running VM
- b. Storage Array Backup: Take a consistent storage snapshot, replicate to tape/VTL media server, remove the snapshot.
Products include Hitachi Shadow Copy, EMC SnapCopy, HP Business/Snap Copy.

Architecture Best Practices for Masking

1. As of release 5.0, the virtualization and masking functions are combined into a single OVA and require additional consideration for installation and configuration. Additionally, support for remote Delphix Masking Engine calls has been implemented and is supported in 5.0.4 and above.
2. Delphix Masking Engines should continue to be deployed to hosts dedicated to that function.
 - Possible exceptions to this would be when any virtualization needed is extremely low and unlikely to be heavily impacted by masking requirements.
3. The Masking Engine is disabled by default.
 - The Delphix Engine will continue to remain running (but unused) on a Masking-only VM.
4. The standard configuration for a dedicated Delphix Masking Engine:
 - a. 8 vCPUs
 - b. 16GB RAM minimum, 32 GB RAM or more recommended.
 - c. 300 GB storage for the OS/system root disk is required for the OS (5.1.4 and greater).
 - d. 50 GB storage for the data disk must be added during the initial configuration via the Engine Setup wizard. (the engine will not complete its setup without a separate data disk).
 - e. If a bulk operation is used, allocate extra space equivalent to the size of all datasets (tables) that will be masked concurrently.
 - i. As a rule of thumb: *Disk Space Required for Bulk = ((Total Database Size * .66) * .10) where Raw data = Total Database Size * .66*
 - ii. 10% change is an estimate based on our experience for data we mask. Often it is lower but there are exceptions such as masking a data warehouse with a large fact table and a bunch of much smaller tables.
 - iii. The VMDK for the engine OS is often stored on the same VMFS volume as the VM definition file (aka VMX). In that case, the VMFS volume must have sufficient space to hold the VMX Configuration, the VMDK for the system disk, and any VMWare logging.
 - iv. Additional VMFS space for swap/paging is required if RAM reservations are not enabled. (The VM will not start if reservations are lacking and disk space is not available for swap)

CPU Utilization

1. One vCPU per concurrent masking job is considered a best practice.
2. Dependent on the algorithms used: Some are calculations such as ones using AES encryption and others are lookups and tend to do more I/O.

Memory Utilization

1. The Delphix Masking Engine uses its memory to cache data. More memory will provide better performance. 1GB per masking job is considered a best practice.
2. Dependent on memory settings in the Masking Engine and JVMs. An increase in parallel workloads will require more memory. Data is either cached directly or using Kettle so the larger the lookups for algorithms the more memory required. This is the first thing to look at for performance issues.

Network and I/O

Delphix Masking leverages the Target DB server and VDB for most of the workload. This means the masking engine can be I/O bound waiting for the DB server. As long as the masking engine can read the data faster than it can process it this is not an issue. Slow networks with numerous hops between the DB server and the Masking server

can cause performance problems. Co-locating the masking server with the DB server is recommended in these cases.

Masking VDB Tuning

1. Always start with the tuning recommendations for Target servers and VDBs first. If the VDB is not performing well, the performance of masking will suffer.
2. For Oracle, it is critical to select no archivelog mode and tune online redo log size at provision time.
3. For SQL Server, the VDB should be in SIMPLE recovery with an appropriate log file and TempDB sizes.

Backup of a Delphix Masking Engine

1. Virtual machine backups are recommended for versions of software in which masking runs in its own VM – in other words, the masking VM is separate from the VM(s) where virtualization takes place.
 - a. If an engine is supporting both masking and virtualization, review data protection best practices.
 - b. Although XML exports of inventories and environments do exist, they are incomplete. Do not rely on them.
2. In-Place (not On-the-Fly) Masking is the primary use case.

The following recommendations apply to the Delphix Masking Engine versions 5.0.2 and earlier:

1. Jobs vs. Streams:
 - a. If there are multiple tables to be masked concurrently, use multiple, separate jobs – one per table.
 - b. Avoid multiple streams due to internal limitations.

 The default setting for streams is 20; set it to 1. This will force serialization (one table at a time) if the job contains multiple tables.

2. Use one update thread per job; this avoids block collisions/contention during the UPDATE phase.

 The default setting is 4; set it to 1.

- a. Identify ALL indexes, constraints, and triggers on columns being masked (and only on columns being masked).
 - b. Evaluate whether it is better to drop/mask/recreate for indexes, or disable/mask/re-enable for triggers and constraints – as compared to leaving in-place during masking. The best choice depends on the situation
3. For Oracle VDBs, use ROWID for SQL UPDATE of masked row value(s).
 - a. Edit the ruleset, select Edit All Logical Keys, enter ROWID as the logical key value; see [Managing Rule Sets](#). When a single, large, non-partitioned table must be masked by concurrent jobs (each masking a subset of the table) to shorten masking elapsed time, segregate jobs by database block/page to avoid contention and locking conflicts. In other words, each job masks a unique set of blocks; each block is masked exclusively by one job.
 - b. If a single table is being masked by multiple, concurrent Jobs, and Indexes/constraints/triggers must be dropped/recreated or disabled/enabled, these must be performed OUTSIDE of masking Jobs.
 - c. Pre-masking and post-masking steps must be created manually.
 - d. Scheduling of pre-script and post-script jobs must be devised. Plan to scheduled/execute externally.

Architecture Best Practices for Source DB and OS Settings

Oracle

1. ARCHIVELOG must be enabled: `select log_mode from v$database.`

2. FORCE LOGGING should be enabled to ensure VDBs are not missing data. When NOLOGGING redo is applied during provision, the resulting VDB will be missing changes. Tables with NOLOGGING changes will throw corruption errors when scanned.
3. Block Change Tracking should be enabled to minimize snapsync time.
4. Consult the documentation for Oracle Standby sources.
5. If the database is encrypted with Oracle TDE (Transparent Data Encryption) plan your Delphix storage requirements with the expectation of minimal compression. Customer Observation: space usage for a TDE dSource copy was 92% (2.44 TB) of the source database size (2.67 TB). A typical Oracle dSource copy for a non-TDE database consumes 40% of the source database size.

SQL Server

1. FULL vs SIMPLE recovery mode trade-off.
2. The maximum size of an MSSQL database that can be linked is 256TB for Windows versions greater than 2003.
The limit is 2TB for Windows 2003. (See [Linking a SQL Server dSource](#))

Architecture Best Practices for Target DB and OS Settings

Target Database Application Settings

1. Oracle:
 - a. Provision with 3 x 5GB online redo logs (minimum) to avoid pause when transaction logs wraparound.
 - b. A provision in NOARCHIVELOG mode to reduce transaction log IO. Masking, Test, QA VDBs rarely need point-in-time rewind
 - c. Always check initialization parameters inherited from a parent, remove any expensive or irrelevant parameters.
 - i. `DB_CACHE_SIZE`, `SGA_TARGET` : set based on the target system being compared to.
 - ii. `FILESYSTEMIO_OPTIONS` to `SETALL` . Any other setting inherited from the source is probably wrong.
 - iii. `DB_BLOCK_CHECKSUM`, `DB_BLOCK_CHECKING`, `DB_LOST_WRITE_PROTECT`, `DB_ULTRA_SAFE` : set to default values to minimize the impact.
 - iv. `PARALLEL_DEGREE_POLICY` to `AUTO` , `PARALLEL_MAX_SERVERS` default, `PARALLEL_EXECUTION_MESSAGE_SIZE` to 32768 (maximum): improve PQ performance.
 - v. `FAST_START_MTTR_TARGET` : drives steady write activity. Set based on the target system being compared to.
 - vi. Consider non-durable commits for Masking, Test, QA, UAT: set `COMMIT_WAIT = NOWAIT` , `COMMIT_LOGGING = BATCH`
 - d. Use Oracle Direct NFS (dNFS) for 11.2.0.4+ (**unstable** on older releases):
 - i. Recommended documentation:
 - [Configuration Examples and Troubleshooting blog](#) from Helmut Hutzler
 - Sample oranfstab to leverage multiple network paths for Delphix VDB
 - ii. Set `DNFS_BATCH_SIZE = 128` (default is 4096). This is a good starting point and sufficient for most workloads.
 - iii. Tune TCP stack: set `tcp_adv_win_scale = 2` due to workaround hard-coded Oracle dNFS TCP buffer size.

- iv. Check Alert Log, V\$DNFS_SERVERS, V\$DNFS_FILES, V\$DNFS_STATS to verify proper working (sample [here](#)).
- e. Create AWR snapshots around a reference customer workload, generate an AWR report.
 - i. AWR snap before/after workload: `SQL> exec DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT();`
 - ii. AWR report between the snaps: `SQL> @?/rdbms/admin/awrrpt`
- f. Generate ASH report to diagnose bottlenecks while a workload is running.
 - `SQL> @?/rdbms/admin/ashrpt`
- g. Run synthetic benchmark `sc-workload`.
- h. Where `db file scattered read` (multiblock cached read) latency is high consult this Support KB: [How to Mitigate Multi-Block Read Performance on Oracle 10g](#)
- i. Improve distributed query performance by modifying dblinks to use local IPs instead of SCAN IPs.
- j. NFS recommended mount options for Oracle RAC/SI: [Oracle Support Note 359515.1](#).

Target Host OS Settings

1. Existing documentation on Target OS practices: [Target Host OS and Database Configuration Options](#)
2. HP-UX 11.31+
 - [Async NFS Direct I/O](#): HP-UX requires Oracle `disk_asynch_io` turned off for filesystems
3. IBM AIX:
 - i. Consult IBM documentation on AIX TCP Tuning
 - ii. [AIX TCP Tuning Prezo](#)
4. Windows:
 - a. Anti-virus programs can impact both performance and operation. Delphix recommends anti-virus scanning exclude folders where Delphix files are maintained, in addition to the normal exclusions put in place for MSSQL operation.
 - b. Delphix Connector (aka DX Connector):
 - i. Plan 3-5GB for the Delphix Connector installation.
 - ii. Windows does not yet have ssh, so Delphix developed the "DX Connector for Windows target host communication.
 - iii. The connector must be installed on all Target Windows hosts.
 - iv. The connector supports two modes – v1 and v2 both use the same application binaries.
 - v. The connector v1 process is used to bootstrap the v2 process on a target. This opens a DSP session back to the Delphix Engine (The same thing is done via SSH on U*nix Targets)
 - vi. v2 mode is required to enable SQL hooks
 - vii. The connector can always be downloaded from a local Delphix Engine at: `http://<delphix_engine>/connector/DelphixConnectorInstaller.msi`.
 - viii. The connector is backward compatible, so it is not always necessary to upgrade it during a Delphix upgrade.
 - c. iSCSI connections:
 - i. Read the following for general awareness of iSCSI limits
 - ii. In addition to the hard limits on iSCSI connections, consideration must be given to the RAM, CPU, and Network to provide sufficient resources for the load on any Target or Staging host.
 - iii. To increase the iSCSI timeout on both Target and Staging hosts.
 - iv. In certain circumstances, it's possible that iSCSI startup will not complete before the SQL Service attempts to start a database. In such circumstances, it can be helpful to ensure the SQL service depends on the iSCSI service.
 - Example: `c:\> sc config "MSSQLServer" depend="Microsoft iSCSI Initiator Service"`
 - v. **Note** that any changes to iSCSI are system-wide and could potentially impact other applications also leveraging that feature.

- d. [Enable Receive Side Scaling \(RSS\)](#) on each network interface that Delphix will be connecting to.

Architecture Best Practices for Staging Targets

This host is called a "staging target" because it has much in common with other targets, such as the remote storage mount to the Delphix Virtualization Engine. Because it leverages remote storage over the network, the staging target only needs enough disk capacity for the OS, database application, and any relevant logs or tools. A staging target is a requirement on all platforms that Delphix supports except Oracle, but you can also use it for Oracle to replay logs leveraging [validated sync](#).

Memory and CPU

1. 32 GB RAM minimum
2. 4 vCPU minimum

General Guidance for Staging Servers (Multi-platform)

1. Delphix recommends dedicated Staging servers for role/architecture separation. However, any Target server can be used as Staging.
2. In cases where the same server is used as both Staging and Target, we strongly recommend a dedicated instance/install for staging to avoid confusion.
3. Delphix recommends at least one Staging server per Delphix Engine to avoid the possibility of a single point of failure across multiple engines.
4. If a staging server is shared among multiple Delphix Engines, please ensure that a dedicated SQL Server Instance is created for each Delphix Engine.
 - Configuration / performance factors:
 - a. Transaction log generation rate.
 - b. Number of VDBs.

 Precise guidance on these items has not yet been defined. In general, if there is a heavy log generation rate and few VDBs, the ideal recommendation is to have at least 1 Staging Target per Delphix Engine.

Disk / Local storage

1. The only local storage needed is for the OS and application with default databases.
2. Storage for a staging database is provided from the Delphix Engine, which is mounted over the network similar to any Target host (NFS/iSCSI).
3. If the customer has a standard DB server build, their standard storage sizing is probably fine.
4. If a recommendation is still needed, suggest 30GB for OS and application and any tools needed.

Network requirements

1. The Staging Target engages in network data transfers between Staging and the Source backup shared location as well as between Staging and the Delphix Engine.
2. The Staging Target is also a Target server, and as such should have < 1ms latency to the Delphix Engine (and low latency to the Source backup, when possible).
3. If the change rate on the Source database(s) is > 1 Gb/sec, the recommended network bandwidth to support network transfers is 10 Gb/sec.
4. In cases where only 1 Gb/sec network bandwidth is available, segregation of each network is recommended to reduce network contention.

5. Ensure that the virtual NIC is using the standard vmxnet3 adapter and not Intel for VMWare based clients. Logical IO errors have been reported while using Intel instead of vmxnet3 adapter.

Windows and MSSQL Specific

1. An MSSQL Server Instance used for Staging should not be clustered.
2. Staging should not be hosted on Windows 2003 - extended support ended July 14, 2015. It is also the first Windows version with iSCSI support and is not ideal.
3. The SQL Server Instances hosted on the Staging Target should have a Maximum Memory set. Also ensure that at all times, at least 10% of total memory is available for OS operations.
4. Only system databases (Master/MSDB/Temp/MSDB) are kept on local storage. All other data is read/written to the Delphix Engine.
5. Windows iSCSI configuration and limits for v2p, target, and staging hosts.
6. [Ensure that Receive Side Scaling \(RSS\)](#) is enabled on every network interface that Delphix will be connecting to.

Related Topics

- [Architecture Checklist FAQ](#)
- [Best Practices](#)

Best Practices for Hypervisor Host and VM Guest

Hypervisor

1. ESXi 6.x or 7.0 is recommended. ESXi 5.5 or earlier is no longer supported.
2. HyperThreading (HT) for Intel®-based servers (no HT on AMD CPUs).
 - a. Disable HT in BIOS, on the ESXi Host, **and** disable [HT Sharing](#) on the Delphix VM for consistency. This is our best practice, disable at all levels.
Any other combination may result in a non-deterministic performance.
 - b. When HT cannot be turned off for both the Host and Delphix VM, it should be turned on at all levels, not run in a "mixed-mode".
 - i. HT disablement at a guest level only can result in non-deterministic performance.
 - ii. A dedicated ESXi host, cluster or DRA is recommended where consistent VDB performance is paramount.
 - c. VM migration to a new host (e.g. VMware HA or vMotion) can create mismatched HT settings.
3. ESXi overhead (resources required for hypervisor cannot be reserved, they must be left unallocated).
 - a. Memory Overhead - 10% of available RAM must not be allocated to guest VMs.
 - *Example: 256GB RAM, allocate 230GB to Delphix VM, leave 26GB for ESX*
 - b. CPU Overhead - At least 2 cores (ideally 4) must not be allocated to guest VMs.
 - *Example: If 16 physical cores are available, allocate 12 to the virtual machines, leaving 4 for the hypervisor*
 - **Why 4?** Certain hypervisor functions require precedence over any virtualized system. If a hypervisor needs more CPU than the amount currently available, it can de-schedule all other virtual processes to ensure adequate CPU resources for the hypervisor. Ensuring the hypervisor will not have to de-schedule any running virtual processes (worlds) by setting aside and not over-subscribing CPUs for virtual functions will leave them available for hypervisor use.
 - c. Even if the Delphix VM is the only VM on a host, the hypervisor is still active and essential; and still needs resources.
4. BIOS Power Management should be set to High Performance where ESXi controls power management.
 - a. It can be impacted by [VMware KB 1018206](#) - poor VM application performance caused by power management settings.
 - b. Ensure that all BIOS managed C-States other than C0 are disabled if power management is hardware controlled.
 - c. Ensure that all ACPI sleep states above S0 are disabled in the BIOS.
 - d. Examples for popular server lines from Cisco, HP, Dell below. *Specific models will vary, use the appropriate spec sheet.*
 - i. **UCS:** disable the Processor Power States, disable Power Technology, set Energy Performance to "Performance"
 - ii. **HP Proliant:** set HP Power Regulator to HP "Static High Performance" mode
 - iii. **Dell:** set BIOS System Profile to "Performance Optimized" mode
5. VMware HA can be enabled; VMware DRS is generally disabled.
6. Blade/Rack Server Firmware and ESXi Drivers should be updated to the latest versions.
7. For Intel®-based servers with E5-2600 v2 processors.
8. Two typical server configurations:
 1. Blade Farm
 2. Rack Server

Hyper-Converged configurations are possible for high performance.

Virtual Machine Guest

1. For VM machine settings, see [Virtual Machine Requirements for VMware Platform](#).
2. VMWare Guest Specifications:
 - a. Minimum: 8 vCPU x 64 GB
Small: 8 vCPU x 128GB
Medium: 16 vCPU x 256 GB
Large: 24 vCPU x 512 GB
 - b. Reserve 100% of RAM and CPU:
 - If the ESX host is dedicated to Delphix, CPU and RAM reservations are advised but not necessary, however, swap space will be required on the hypervisor to compensate for the lack of reserved RAM.
 - c. [Hyperthreading](#) - See the ESX host section at the top. Disable HT Sharing on VM, disable HT on ESX Host.
3. Assign single-core sockets for vCPUs in all cases. If there is a compelling reason to use multi-core vCPUs, reference the following article from VMware which describes matching virtual multi-core sockets to the hardware ESX is running on.
[VMware Article on CoresPerSocket](#)
Example: *ESXi Host has 2 socket x 18 core Intel Xeon, Delphix Engine wants 16 vCPU.*
Configure Delphix VM with 2 Virtual Sockets, 8 Cores Per Socket to utilize hardware architecture.
4. Avoid placing other extremely active VMs on the same ESX host.
5. Monitoring - vSphere Threshold Alerts for CPU, Network, Memory, Capacity.
6. To set the number of vCPUs per virtual machine via the vSphere client, please see "[Virtual CPU Configuration](#)" in the Administration guide:
[ESXi 6.0](#)
 - a. [Delphix VM CPU Utilization](#) - Delphix KB article on what makes Delphix VMs similar to other resource-intensive applications
 - b. [Exchange on VMware Best Practices](#) - VMworld 2013 session
 - c. [ESXTOP Reference, Blog](#)
7. Ensure that the latest available VMware drivers and firmware versions are installed for HBAs, NICs and any other hardware components configured on the Delphix virtual machine. This is a critical step that can have a massive impact on the performance and robustness of our solution.

Related Topics

- [Best Practices Hypervisor Host and VM Guest ESX FAQ](#)
- [Best Practices](#)

Best Practices Hypervisor Host and VM Guest ESX FAQ

Frequently Asked Questions

Why does Delphix require a minimum of 8 vCPUs and recommend 128 GB per 8 vCPUs?

8 vCPUs are not only a standard licensing block, but they are also key to meeting our 10Gbps single-engine throughput potential and help to sustain low latency for VDBs.

As with CPU, cache memory is required to drive peak loads on the Delphix Engine. More memory allows for more blocks to be read from the cache rather than going to less performant disks. Delphix stores cached data in a compressed format and only keeps a single copy of unique blocks in memory. These features give read performance across multiple VDBs provisioned from a single source dramatic improvements in speed, scalability and memory utilization.

Why does Delphix request reservations for CPU and memory?

Delphix performance can be greatly impacted there is contention over CPU or RAM. Reservations allow the engine to explicitly control those resources and avoids the possibility of contention with other VMs, even when resources are overcommitted.

Why does Delphix request Hyper-Threading (HT) be disabled?

Hyper-Threading can have a positive impact on many, but not all applications. Delphix has a different execution profile that does not benefit from Hyper-Threading. Factors such as significant memory bandwidth requirements and a high level of parallelism which requires a high number of shared locks mean that HT and Delphix do not generally work well together.

Why does Delphix request 4 controllers, and why must the storage be identical between them?

To provide optimal storage performance, you must spread data equally over the maximum (4) virtual SCSI controllers. To provide consistent performance between each of the four controllers, you need to ensure storage is identical between them.

Why must virtual disks (VMDKs) be thick provisioned and eager-zeroed?

Thick provisioning and eager zero ensure performance is top-notch from the start with no hiccups from expanding virtual resources.

Why is 20% free space required?

While the ZFS file system has a lot of features leveraged by Delphix, it loses efficiency as space decreases. 20% is the minimum that must be available for best performance.

Why does Delphix request you reserve CPU and RAM for Hypervisor overhead?

Based on VMWare's resource management guide and our own experiences with high IO throughput. Note there is no specific mechanism to assign resources to the hypervisor, the only way to preserve overhead is by not allocating resources to guests.

Why does Delphix generally want VMWare HA enabled, but DRS disabled?

VMware HA (High Availability) addresses outages that occur when a physical host goes down or is completely offline, by migrating the guest(s) to another physical host and restarting them. There is no real downside, it simply brings unavailable servers back online.

VMware DRS (Distributed Resource Scheduler) is for load balancing host resources in a cluster. Because of high IO and best practices configuration for optimal performance, our engine is typically not a good candidate for relocation.

Why does Delphix request you set power management to High-Performance Mode?

This will ensure power management will never impact performance by entering into a lower power state (also known as c-state).

Related Topics

- [Best Practices for Hypervisor Host and VM Guest](#)
- [Best Practices](#)

Best Practices for Network Configuration

1. For more information about network configuration refer to [Network Performance Configuration Options](#).
2. Delphix Engine <====> Target Host:

Implement standard requirements for optimal NFS/iSCSI performance:

 - a. Optimal physical network topology:
 - Low latency: < 1ms for 8K packets.
 - Network adjacency: minimize network hops, co-locate in the same blade enclosure, co-locate on the same physical host.
 - Eliminate all Layer 3+ devices - firewalls, IDS, packet filters (Deep Packet Inspection - DPI).
 - Multiple switches can add latency and fragmentation and reordering issues will add significant latency.
 - b. Optimal throughput:
 - 10GbE physical uplinks or higher.
 - Jumbo frames (typically MTU 9000) improves network efficiency for 8K packets: lower CPU, lower latency, greater throughput.
 - All devices end-to-end *must* be configured for the larger frame size including switches, routers, fabric interconnects, hypervisors and servers.
 - Traffic between two VMs on the same host is limited to ~16Gbps when leveraging built-in virtual networking.
 - c. Optimal logical flow:
 - Disable QoS throttles limiting network utilization below line rate (e.g. HP Virtual Connect FlexFabric).
 - Consider a dedicated VLAN (with jumbo frames) for NFS/iSCSI traffic.
 - [VMware KB-Configuring iSCSI port binding with multiple NICs in one vSwitch for VMware ESXi 6.0.x \(2045040\)](#)
 - [VMware KB-Considerations for using software iSCSI port binding in ESX/ESXi \(2038869\)](#)
 - d. NIC Teaming (at ESX layer) of multiple physical uplinks can provide additional throughput for higher workloads.
 - i. Examples: 4 x 1Gb NICs support up to 400 MBPS IO, 2 x 10Gb NICs support up to 2 GBPS IO.
 - ii. [VMware KB-1004088](#) has NIC teaming recommendations, including `route-based-on-IP-hash` policy.
 - iii. [VMware KB-1001938](#) has host requirements for physical link aggregation (LACP, EtherChannel).
 - iv. [VMware KB-1007371](#), [popular blog post](#) details problems with NIC selection using `dest-IP hash`.
 - e. Fragmentation and dropped packets can result in excessive retransmissions of data, reducing throughput.
 - i. On AIX, LSO (Large Send Offload) and LRO (Large Receive Offload) network features have caused many problems. The virtualization engine employs LSO only, but no LRO, AIX employs both. These features are enabled by default. Best practice is to disable LSO on the VE when the target is AIX.
 - LSO enabled on the Virtualization engine can result in dramatically limited or blocked transmit throughput (and increase re-transmission traffic).
 - When wishing to disable all, LSO should be disabled on the Delphix Virtualization Engine, while both LRO and LSO should be disabled on guest VM of AIX Target. The specific commands to disable LSO and LRO on AIX are not captured here and appear to be context/version-specific.
 - f. **Jumbo frames check via ping**

```
Delphix Engine
$ ping -D -s [Target_IP] 8000
"ICMP Fragmentation needed and DF set from gateway" indicates MTU < 8028
```

```
Linux
$ ping -M do -s 8000 [Delphix_Engine_IP]
"Frag needed and DF set (mtu = xxxx)" indicates MTU < 8028
```

```
MacOS
ping -D -s 8000 [Delphix_Engine_IP]
```

Note: "sudo sysctl -w net.inet.raw.maxdgram=16384" will increase the max ICMP datagram size on Mac, allowing you to use -s 9000 on Mac0

```
Windows
ping -f -l 8000 [Delphix_Engine_IP]
```

<http://www.mylesgray.com/hardware/test-jumbo-frames-working/>

- g. Measure Network Bandwidth and Latency:
 - i. Latency in both directions should be < 1ms for an 8KB payload.
 - ii. Network Hops should be minimized: `tracert` (Unix/Linux) / `tracert` (windows).
 - iii. Throughput in both directions: 50-100 MB/s on 1 GbE, 500-1000 MB/s on 10 GbE physical link.
 - iv. Always use the CLI when possible.
- h. NIC should use Auto-negotiate on Ethernet with a minimum of 1000Mbps.
 - Hard setting speed/duplex will limit network throughput below the line rate.
3. **Delphix <====> Staging Server** (SQL Server, Sybase):
Measure latency, bandwidth for transaction log restore performance
4. **Source <====> Delphix:**
Measure latency, bandwidth to explain snapsync performance
5. **ESX host <====> ESX host** (ESX Cluster):
 - i. This is because the entire memory footprint of the Delphix VM (more precisely, the entire address space of the ESX processes that comprise the VM) must be copied to the receiving ESX host, along with all changes to that address space as they happen. Our ZFS cache comprises the bulk of that memory and changes as I/O occurs, and that rate of change is governed by the network bandwidth available to the Delphix Engine.
 - ii. Attempting to live vMotion a DE with a 10Gb NIC over a 1Gb vMotion network is thus likely to fail.

Related Topics

- [Best Practices Network FAQ](#)
- [Best Practices](#)

Best Practices Network FAQ

Frequently Asked Questions

Why does Delphix request 10GE Ethernet?

As a matter of physics and standards - 10 gigabit (Gb) Ethernet can sustain approximately 1 gigabyte (GB) per second of throughput. With all our best practices applied, a Delphix Engine can achieve very close to that line speed, allowing for optimal load, engine, and license utilization. Lower network speeds may be acceptable for low loads, while in some environments NIC teaming (e.g. LACP) may be required for top speeds.

Why does Delphix require < 1ms latency to TARGET servers and < 50ms to SOURCE servers?

Delphix leverages NFS and iSCSI (depending on platform) for live TARGET DB mounting over the network, so it's imperative that latency is as low as possible. Data coming from SOURCE servers is not generally as time-sensitive, so you need a minimum latency of < 50ms to ensure operational integrity.

Why does Delphix request Jumbo Frames?

Jumbo frames increase the Ethernet maximum transmission unit (MTU) from the default 1500 bytes to 9000 bytes. This has several effects such as decreasing CPU cycles by transferring fewer packets and increasing the engine throughput. You will find jumbo frames have a 10-20% real-world impact and are required (along with all other best practices) to handle peak loads of 800 - 1000MB/s on an 8 vCPU engine with a 10Gb network.

How does Delphix avoid communication impact with non-jumbo frame hosts when Jumbo Frames are enabled on the Delphix Host?

[Path MTU Discovery](#) is the mechanism by which two hosts agree on the MTU leveraged for communication between them. This mechanism will ensure communication between both standard and Jumbo Frame enabled hosts works as expected.

When does Delphix recommend NIC teaming?

The Delphix Engine is capable of high throughput, but not every enterprise has sufficient network bandwidth to support it. Teaming is a less expensive way of increasing the bandwidth when compared to new hardware.

Why does Delphix recommend logical and physical and co-location?

The Delphix Engine leverages network connections extensively, so optimizing the latency whenever possible is very important - sometimes critical.

Related Topics

- [Best Practices for Network Configuration](#)
- [Best Practices](#)

Best Practices for Storage

1. Configuration:

- a. Virtual Disks (VMDK) with spinning or tiered media must be thick provisioned + lazy zeroed.
 - For storage which is 100% SSD/EFD/Flash-based, continue to thick provision, however, eager zero is not necessary.
- b. VMDKs may be homed on VMFS or RDM storage, however, VMFS is much more common and generally preferred.
- c. Regardless of VMFS/RDM selection, physical LUNs must have uniform characteristics (RAID, spindle count, tier) and should be thin provisioned to save time.
- d. Storage allocated must be identical between each vSCSI controller.
- e. The supported maximum of four virtual SCSI controllers (PVSCSI (default) or LSI Logic Parallel) must be enabled. A mix of different types of SCSI controllers is not supported within the engine.
- f. Virtual Disks must be evenly distributed across the 4 virtual SCSI controllers. For example, 8 virtual disks should be configured as 2 disks per controller: SCSI (0:0), SCSI (0:1), SCSI (1:0), SCSI (1:1), SCSI (2:0), SCSI (2:1), SCSI (3:0), SCSI (3:1).
- g.
 - You don't need to account for the OS in the even distribution of disks across controllers, just pick one; the OS doesn't place a substantial load on the controller.
- h. To provision VMDK disks over 16TB in size, the vSphere web client must be used, the win32 client will return an error.
- i. Delphix requires 127GB for the system disk. older versions defaulted to 150GB.
 - Due to our unified OVA for masking, the 127GB requirement also applies to Masking Engines
- j. VMDK for the Delphix Engine OS is often stored on the same VMFS volume as the Delphix VM definition file (aka VMX). In that case, the VMFS volume must have sufficient space to hold the Delphix VMX Configuration, the VDMK for the system disk, a swap/paging area if the memory reservation was not enabled for the Delphix Engine (or to suspend the system), and any related VMware logging.
- k. Set ESX Storage Multipathing IO optimization ([KB 2072070](#)).
 - i. Set multipathing policy to round-robin.
 - ii. Set path switching IO operation limit to 1 (default 1000).
- l. Verify that the queue depth setting on the virtual disks used for VMware is appropriate based on the minimum of the HBA type or the underlying storage's combined queue depth for all hosts attached to the same controller
 - See this [VMware KB article](#) for how to check/set the queue depth

2. Testing:

- a. Run Storage Performance Tool on the raw storage **before any engine configuration**. This is a one-time opportunity for each engine upon installation
- b. Required maximum storage latency is < 2ms for writes and < 10ms (95th percentile) for small random reads.
Minimum passing grades: 4KB/8KB reads (B-), 1MB reads (A), 1KB/128KB writes (A).
- c. If working with the Delphix Professional Services team, we would expect to run additional baseline performance measurements via composite tools and scripts.
 - e.g. "Sanity Check" (Oracle) or "DiskSpd" (MSSQL)

3. Detail Discussion:

- a. Before beginning any discussion on storage performance, or at the beginning of the installation, collecting the following specs from your storage administrator will assist in understanding
- b. Vendor, Model (For example: EMC VMAX 40k, HP 3PAR StoreServ 7000)
- c. IO latency SLO (For example 5ms 99.999%)
- d. IOPS/GB SLO (For example: 0.68 IOPS/GB for EMC Gold-1)
- e. Cache type and size (For example FAST cache 768GB)
- f. Tier, #Pools; if auto – tiering; relocation schedule (For example: Gold/Silver/Auto/3 pools/etc)
- g. Pool detail: (#) drives, RPM, Type (For example: Pool1: (20) EFD, (30) 15k SAS, Pool 2: (40) 10k SATA)

- h. Connection (For example: 16Gb Fibre Channel, 10Gb iSCSI, 20Gb NFS)
- i. Dedicated or Shared pool (how many applications/servers)

Related Topics

- [Best Practices](#)

Best Practices for Delphix Engine Data Protection

The following are recommended best practices for data protection in the Delphix Engine.

1. For protection against physical host failure - leverage VMware HA.
2. For protection against storage failure - leverage Delphix replication.
3. For protection against administrative error - leverage storage snapshots.
4. For protection against site failure - leverage Delphix replication and/or Delphix Live Archive.
5. Infrastructure Backup of the Delphix VM:
Must take a consistent group snapshot of all Delphix VM storage (system disk, VM configuration, database VMDK/RDMs)
RTO, RPO are inferior compared to Delphix Replication. The granularity of restore is at the VM level: **all-or-nothing**.
 - a. Virtual Machine Backup: Create VM snapshot, backup (proxy server), remove the snapshot.
Products use VMware APIs: NBU for VMware, TSM, NetWorker
Limited to < 2TB because of time to backup, impact on running VM
 - b. Storage Array Backup: Take a consistent storage snapshot, replicate to tape/VTL media server, remove the snapshot.
Products include Hitachi Shadow Copy, EMC SnapCopy, HP Business/Snap Copy.

Related Topics

- [Best Practices Delphix Engine Data Protection FAQ](#)
- [Best Practices](#)

Best Practices Delphix Engine Data Protection FAQ

Frequently Asked Questions

Why does Delphix recommend SAN snapshots or Delphix replication for backup?

There are a few possible methods for data protection of the Delphix Engine. Those methods are SAN snapshots, Delphix replication, and virtual machine snapshots (for very small engines only). Because the Delphix Engine is itself a backup of source environments, many customers simply plan to rebuild in the event of a disaster.

What is the DXToolkit, and how can it help?

The professional services team has created a Perl-based “DXToolkit” which can help export and import certain configuration data over web services. This toolkit can be leveraged to assist with would normally be a manual re-install outside of the above methods.

For further detail around data protection, please speak with your Delphix contact.

Can I use a VMware-based backup solution such as VEEAM to backup my Delphix Engine?

Yes, VMWare backup solutions are useful for backing up guest VMs. However, Delphix suggests that you only use this approach for Delphix Engines which have a smaller storage footprint (perhaps < 2 TB) and are less active.

Running this type of backup puts a load on the environment, which might adversely impact Delphix VM performance.

Can I use a VMware snapshot for backing up Delphix for a small window – for example, during an engine upgrade?

Yes. However, even though snapshots are instantaneous, they track changes separately from the base disks and can grow to consume as much space as the original.

 Upgrades, in particular, can change substantial amounts of data. If you lose physical disks, snapshots are useless because it needs them to make up the current state of VM. A Delphix Engine is often allocated multiple terabytes of storage and is often very busy due to load aggregation from virtual databases on multiple target servers, so this approach may be challenging. Snapshots cannot detect storage corruption.

Can I use a Storage snapshot solution to protect Delphix against Storage and Delphix corruption?

Yes. However, please note that the caveats which apply to VMWare snapshots will also apply here.

 A specific concern related to storage layer snapshots is that you must create a consistency group that contains both the OS and Data disks.

Can I use RMAN to backup my VDBs just like a Physical database to provide extra protection?

You can backup Delphix VDBs using Oracle RMAN tools, but the recovery database would first require re-hydration of that VDB, which might take up equivalent production storage space.

Furthermore, that re-hydrated database needs to be brought into the Delphix framework as a dSource, after which you can provision a VDB to complete recovery. The whole process might take hours or days to recover.

The best approach is to use the VDB Snapshot capability to backup VDB frequently and then leverage Delphix Replication capability to protect underlying Delphix storage, which holds that VDB snapshot.

Related Topics

- [Best Practices for Delphix Engine Data Protection](#)
- [Best Practices](#)

Best Practices for Source DB and OS Settings

Best Practices

Oracle

1. ARCHIVELOG must be enabled: `select log_mode from v$database.`
2. FORCE LOGGING should be enabled to ensure VDBs are not missing data. When NOLOGGING redo is applied during provision, the resulting VDB will be missing changes. Tables with NOLOGGING changes will throw corruption errors when scanned.
3. Block Change Tracking should be enabled to minimize snapsync time.
4. Consult the documentation for Oracle Standby sources.
5. If the database is encrypted with Oracle TDE (Transparent Data Encryption) plan your Delphix storage requirements with the expectation of minimal compression. Customer Observation: space usage for a TDE dSource copy was 92% (2.44 TB) of the source database size (2.67 TB). A typical Oracle dSource copy for a non-TDE database consumes 40% of the source database size.

SQL Server

1. If using a FULL recovery model, configure your dSource to stay synchronized using Transaction Log backups. This will usually allow the dSource to stay in sync using much less network and disk IO than Full or Differential backups.
2. Ensure that the number of Virtual Log Files (VLFs) in the Source database is appropriate. Databases with hundreds or thousands of VLFs will experience slower provisioning times as SQL Server must do more work during recovery. The blog post [A Busy/Accidental DBA's Guide to Managing VLFs](#) has helpful information here.
3. Review index maintenance and defragmentation operations on your Source database, to ensure that they are not running more often than necessary. Many maintenance operations are logged, and can result in extremely large log backups without much benefit to the database server. See [Stop Worrying About SQL Server Fragmentation](#).
4. If configuring a dSource to use Full or Differential backups, ensure that Source database backups do not run at the same time as database maintenance or large batch operations. This can significantly increase the amount of time required to perform database recovery and will slow down the provisioning of VDBs.
5. For best compatibility with the Delphix Engine, avoid taking log backups more frequently than every 10 minutes. The validated sync process requires time to detect and validate log backups before they can be applied to a dSource, and extremely high-frequency log backups can make it difficult for a dSource to stay in sync.
6. If you are using SQL Authentication for database logins, and you want Source database users to be mapped to database logins when provisioning VDBs, ensure that the SID of your database logins is the same between your Source and Target environments (the passwords can be different). For more details, see the [Correcting Orphaned SQL Server Database Users \(KBA1111\)](#) KB article.
7. To maximize dSource performance, ensure that backups are being taken to a disk drive or network location that can be accessed by the Staging Server at high speed.

Related Topics

- [Best Practices](#)
- [Oracle Environments and Data Sources](#)
- [SQL Server Environments and Data Sources](#)

Best Practices for Target DB and OS Settings

Target Database Application Settings

1. Oracle:
 - a. Provision with 3 x 5GB online redo logs (minimum) to avoid pause when transaction logs wraparound.
 - b. Provision in NOARCHIVELOG mode to reduce transaction log IO. Masking, Test, QA VDBs rarely need point-in-time rewind
 - c. Always check initialization parameters inherited from parent, remove any expensive or irrelevant parameters.
 - i. `DB_CACHE_SIZE`, `SGA_TARGET` : set based on target system being compared to.
 - ii. `FILESYSTEMIO_OPTIONS` to `SETALL` . Any other setting inherited from source is probably wrong.
 - iii. `DB_BLOCK_CHECKSUM`, `DB_BLOCK_CHECKING`, `DB_LOST_WRITE_PROTECT`, `DB_ULTRA_SAFE` : set to default values to minimize impact.
 - iv. `PARALLEL_DEGREE_POLICY` to `AUTO` , `PARALLEL_MAX_SERVERS` default, `PARALLEL_EXECUTION_MESSAGE_SIZE` to 32768 (maximum): improve PQ performance.
 - v. `FAST_START_MTTR_TARGET` : drives steady write activity. Set based on target system being compared to.
 - vi. Consider non-durable commits for Masking, Test, QA, UAT: set `COMMIT_WAIT = NOWAIT`, `COMMIT_LOGGING = BATCH`
 - d. Use Oracle Direct NFS (dNFS) for 11.2.0.4+ (**unstable** on older releases):
 - i. Recommended documentation:
 - [Oracle Direct NFS Configuration](#)
 - [Configuration Examples and Troubleshooting blog](#) from Helmut Hutzler
 - Sample oranfstab to leverage multiple network paths for Delphix VDB
 - ii. Set `DNFS_BATCH_SIZE = 128` (default is 4096). This is a good starting point and sufficient for most workloads.
 - iii. Tune TCP stack: set `tcp_adv_win_scale = 2` due to workaround hard-coded Oracle dNFS TCP buffer size.
 - iv. Check Alert Log, `V$DNFS_SERVERS`, `V$DNFS_FILES`, `V$DNFS_STATS` to verify proper working (sample [here](#)).
 - e. Create AWR snapshots around a reference customer workload, generate an AWR report.
 - i. AWR snap before/after workload: `SQL> exec DBMS_WORKLOAD_REPOSITORY.CREATE_SNAPSHOT();`
 - ii. AWR report between the snaps: `SQL> @?/rdbms/admin/awrrpte`
 - f. Generate ASH report to diagnose bottlenecks while a workload is running.
 - `SQL> @?/rdbms/admin/ashrpt`
 - g. Run synthetic benchmark `sc-workload`.
 - h. Where `db file scattered read` (multiblock cached read) latency is high consult this Support KB: [How to Mitigate Multi-Block Read Performance on Oracle 10g](#)
 - i. Improve distributed query performance by modifying dblinks to use local IPs instead of SCAN IPs.

- j. NFS recommended mount options for Oracle RAC/SI: [Oracle Support Note 359515.1](#).

Target Host OS Settings

1. Existing documentation on Target OS practices: [Target Host Configuration Options for Improved Performance](#)
2. HP-UX 11.31+
 - [Async NFS Direct I/O](#): HP-UX requires Oracle `disk_asynch_io` turned off for filesystems
3. IBM AIX:
 - i. Consult IBM documentation on AIX TCP Tuning
 - ii. [AIX TCP Tuning Prezo](#)
4. Windows:
 - a. Anti-virus programs can impact both performance and operation. Delphix recommends anti-virus scanning exclude folders where Delphix files are maintained, in addition to the normal exclusions put in place for MSSQL operation.
 - b. Delphix Connector (aka DX Connector):
 - i. Plan 3-5GB for the Delphix Connector installation.
 - ii. Windows does not yet have ssh, so Delphix developed the "DX Connector for Windows target host communication.
 - iii. The connector must be installed on all Target Windows hosts.
 - iv. The connector supports two modes – v1 and v2 both use the same application binaries.
 - v. The connector v1 process is used to bootstrap the v2 process on a target. This opens a DSP session back to the Delphix Engine (The same thing is done via SSH on U*nix Targets)
 - vi. v2 mode is required to enable SQL hooks
 - vii. The connector can always be downloaded from a local Delphix Engine at: `http://<delphix_engine>/connector/DelphixConnectorInstaller.exe`.
 - viii. The connector is backwards compatible, so it is not always necessary to upgrade it during a Delphix upgrade.
 - c. iSCSI connections:
 - i. Read the following for general awareness of iSCSI limits
 - ii. In addition to the hard limits on iSCSI connections, consideration must be given to the RAM, CPU and Network to provide sufficient resources for the load on any Target or Staging host.
 - iii. To increase the iSCSI timeout on both Target and Staging hosts.
 - iv. In certain circumstances it's possible that iSCSI startup will not complete before the SQL Service attempts to start a database. In such circumstances, it can be helpful to ensure the SQL service depends on the iSCSI service.
 - Example: `c:\> sc config "MSSQLServer" depend="Microsoft iSCSI Initiator Service"`
 - v. **Note** that any changes to iSCSI are system-wide and could potentially impact other applications also leveraging that feature.
 - d. [Enable Receive Side Scaling \(RSS\)](#) on each network interface that Delphix will be connecting to.

Related Topics

- [Best Practices](#)
- [Architecture Checklist](#)

Receive Side Scaling (RSS) for Windows Staging Target and Targets

Enabling Receive Side Scaling (RSS) on a Windows Target and Staging Target can have a significant improvement in the overall IO throughput to the Delphix Engine and is a best practice. RSS enables network adapters to distribute the kernel-mode network processing load across multiple processor cores in multi-core computers. The distribution of this processing makes it possible to support higher network traffic loads than would be possible if only a single core were to be used.


More information on RSS can be found [here](#).

⚠ Enabling RSS on the network interface will force the network service to restart and will cause a momentary loss of connectivity on that network interface.


i Because hyper-threaded CPUs on the same core processor share the same execution engine, the effect is not the same as having multiple core processors. For this reason, RSS does not use hyper-threaded processors.

Steps to implement RSS on Windows


1. From Server Manager/Local Server/Network Connections select the NIC that Delphix will be connecting to.



2. Select Properties and then Configure.



3. From the Property menu on the left, select Receive Side Scaling and select 'OK' to close each of the open windows.



Related Topics

- [Best Practices](#)
- [Architecture Checklist](#)

Best Practices for Staging Targets

This host is called a "staging target" because it has much in common with other targets, such as the remote storage mount to the Delphix Virtualization Engine. Because it leverages remote storage over the network, the staging target only needs enough disk capacity for the OS, database application, and any relevant logs or tools. A staging target is a requirement on all platforms that Delphix supports except Oracle, but you can also use it for Oracle to replay logs leveraging [validated sync](#).

Memory and CPU

1. 32 GB RAM minimum
2. 4 vCPU minimum

General Guidance for Staging Servers (Multi-platform)

1. Delphix recommends dedicated Staging servers for role/architecture separation. However, any Target server can be used as Staging.
2. In cases where the same server is used as both Staging and Target, we strongly recommend a dedicated instance/install for staging to avoid confusion.
3. Delphix recommends at least one Staging server per Delphix Engine to avoid the possibility of a single point of failure across multiple engines.
4. If a staging server is shared among multiple Delphix Engines, please ensure that a dedicated SQL Server Instance is created for each Delphix Engine.

Configuration / performance factors:

- a. Transaction log generation rate.
- b. Number of VDBs.

 Precise guidance on these items has not yet been defined. In general, if there is a heavy log generation rate and few VDBs, the ideal recommendation is to have at least 1 Staging Target per Delphix Engine.

Disk / Local Storage

1. The only local storage needed is for the OS and application with default databases.
2. Storage for a staging database is provided from the Delphix Engine, which is mounted over the network similar to any Target host (NFS/iSCSI).
3. If the customer has a standard DB server build, their standard storage sizing is probably fine.
4. If a recommendation is still needed, suggest 30GB for OS and application and any tools needed.

Network Requirements

1. The Staging Target engages in network data transfers between Staging and the Source backup shared location as well as between Staging and the Delphix Engine.
2. The Staging Target is also a Target server, and as such should have < 1ms latency to the Delphix Engine (and low latency to the Source backup, when possible).
3. If the change rate on the Source database(s) is > 1 Gb/sec, the recommended network bandwidth to support network transfers is 10 Gb/sec.
4. In cases where only 1 Gb/sec network bandwidth is available, segregation of each network is recommended to reduce network contention.

5. Ensure that the virtual NIC is using the standard vmxnet3 adapter and not Intel for VMWare based clients. Logical IO errors have been reported while using Intel instead of vmxnet3 adapter.

Windows and MSSQL Specific

1. An MSSQL Server Instance used for Staging should not be clustered.
2. Staging should not be hosted on Windows 2003 - extended support ended July 14, 2015. It is also the first Windows version with iSCSI support and is not ideal.
3. The SQL Server Instances hosted on the Staging Target should have a Maximum Memory set. Also ensure that at all times, at least 10% of total memory is available for OS operations.
4. Only system databases (Master/MSDB/Temp/MSDB) are kept on local storage. All other data is read/written to the Delphix Engine.
5. Windows iSCSI configuration and limits for v2p, target and staging hosts.
6. [Ensure that Receive Side Scaling \(RSS\)](#) is enabled on every network interface that Delphix will be connecting to.

Related Topics

- [Best Practices](#)
- [Enabling Validated Sync for Oracle](#)
- [Best Practices for Target DB and OS Settings](#)

Data Backup and Recovery Solutions

 Cloning an existing Delphix Engine by using hypervisor VM cloning features or by copying existing Engine storage to a new virtual machine is currently not supported. You can choose from the following methods for backup or duplication of an Engine.

 Unknown macro: 'navbox'

Delphix Data Vault

Overview

The Delphix Engine has a base feature set that is compelling as a data protection solution. An enhancement is being introduced that can further prevent snapshot and database loss in the event of a ransomware attack.

Delphix Data Vault for ransomware protection allows organizations to recover their application data access much faster than traditional backup solutions, in case of malicious attacks. It relies on the new **Data Vault Replication** feature, which replicates critical business DB data stored on Delphix engines to a new target engine called Data Vault. Once securely stored on the Data Vault, the replicated DB data can be used to recover business applications with very low RTO and RPO.

This feature is available via CLI or via the Data Vault UI section added to the Replication page in versions 6.0.9.0 and above. The UI provides function for creating Data Vault profiles from scratch and convert existing specifications to be locked.

Requires Technical Services Consult

Delphix **requires** a Technical Services assessment prior to deployment and configuration of Data Vault. The process of configuring a Delphix Data Vault replication profile is simple; the assessment is required because each application has specific data protection and recovery requirements and we must ensure that Delphix can respond to them appropriately. To schedule an assessment, please contact your Customer Success Manager.

Advantages

The Delphix Data Vault solution provides built-in protection against ransomware attacks by frequently refreshing production application data to a virtual database (dSource), from which you can recover applications almost instantly to any valid point in time prior to the encryption attack by leveraging Delphix VDB provisioning capabilities. The refresh interval between the production application and the Delphix dSource can be configured to be up to seconds, for business-critical applications (not applicable to all DB technologies).

- Creates a separation of responsibilities between the two Delphix Engines
- One engine is used for regular Virtualization cases (ingestion, VDBs, SDD)
- Another engine is used for ransomware protection.
- Creates a physical separation by allowing the admin to isolate and secure the locked Delphix Engine.
- Only the DSP port has to be open for replication.
- No ports are needed for JDBC, NFS, or SSH until VDBs need to be created to export data. This also prevents attack vectors related to any of those protocols.
- Can assist with making deployments and security reviews easier to pass since the locked Delphix Engine is isolated and has a single purpose to reduce potential attack vectors.

In the event of a ransomware attack on a primary engine source being compromised or corrupted, you can provision a VDB on the replica in the locked namespace of the replication target – similar to the normal replication namespace. This process can be further outlined in the [Provisioning From Replicated Data Sources or VDBs](#) article. If a complete recovery of the primary engine is needed, please contact Delphix Support.

Implementation

This feature adds a property to the replication namespace and specifications called “locked”. Additional dSources, VDBs, groups, and domains can be added to locked replication specs, but data sources cannot be removed after doing so. Failover on the target Delphix Engine is not allowed if the namespace is part of the locked replication spec. The retention policy duration on a locked namespace can be modified as long as the duration is either being increased, or it is being decreased to a minimum of 100 days.

The time configuration on Delphix engines with Data Vault enabled cannot be changed. This is to prevent attempts at bypassing retention policies in order to try and delete snapshots on the target data vault.

A fault is now generated on the Data Vault target for a locked namespace that has not received a successful replication update in 3 days.

CLI Functions

Create a locked replication profile.

```
[user.hostname]> replication spec
[user.hostname] replication spec> create
[user.hostname] replication spec create *> set name=locked-spec-1
[user.hostname] replication spec create *> set objectSpecification.objects=Untitled/
dbname
[user.hostname] replication spec create *> set targetHost=example.delphix.com
[user.hostname] replication spec create *> set targetPrincipal=admin
[user.hostname] replication spec create *> set targetCredential.password=delphix
[user.hostname] replication spec create *> set lockedProfile=true
[user.hostname] replication spec create *> commit
`REPLICATION_SPEC-3
[user.hostname] replication spec> select locked-spec-1
[user.hostname] replication spec 'locked-spec-1'> get
  type: ReplicationSpec
  name: locked-spec-1
  automaticReplication: false
  bandwidthLimit: 0
  description: (unset)
  encrypted: false
  lockedProfile: true <----- LOCKED
  numberOfConnections: 1
  objectSpecification:
    type: ReplicationList
    name: (unset)
    objects: Untitled/dbname
  reference: REPLICATION_SPEC-3
  runtime:
    type: ReplicationSpecRuntime
  schedule: (unset)
  tag: 5570be25-dbcf-48c3-b2d2-dd2c65eb98b7
  targetCredential:
    type: PasswordCredential
    password: *****
  targetHost: example.delphix.com
  targetPort: 8415
```

```

    targetPrincipal: admin
    useSystemSocksSetting: false
[user.hostname] replication spec 'locked-spec-1'> cd ..
[user.hostname] replication spec>

```

Lock an unlocked replication profile.

```

[user.hostname]> replication spec create
[user.hostname] replication spec create *> set name=locked-spec-2
[user.hostname] replication spec create *> set objectSpecification.objects=Untitled/
dbname
[user.hostname] replication spec create *> set targetHost=example.delphix.com
[user.hostname] replication spec create *> set targetPrincipal=admin
[user.hostname] replication spec create *> set targetCredential.password=delphix
[user.hostname] replication spec create *> commit
`REPLICATION_SPEC-4
[user.hostname]> replication spec select locked-spec-2
[user.hostname] replication spec 'locked-spec-2'> get
  type: ReplicationSpec
  name: locked-spec-2
  automaticReplication: false
  bandwidthLimit: 0
  description: (unset)
  encrypted: false
  lockedProfile: false
  numberOfConnections: 1
  objectSpecification:
    type: ReplicationList
    name: (unset)
    objects: Untitled/dbname
  reference: REPLICATION_SPEC-4
  runtime:
    type: ReplicationSpecRuntime
  schedule: (unset)
  tag: e8608d05-0693-440d-8a2b-8c6cbfe06a62
  targetCredential:
    type: PasswordCredential
    password: *****
  targetHost: example.delphix.com
  targetPort: 8415
  targetPrincipal: admin
  useSystemSocksSetting: false
[user.hostname] replication spec 'locked-spec-2'> update
[user.hostname] replication spec 'locked-spec-2' update *> set lockedProfile=true
[user.hostname] replication spec 'locked-spec-2' update *> commit
[user.hostname] replication spec 'locked-spec-2'> get lockedProfile
true
[user.hostname] replication spec 'locked-spec-2'>

```

Verify the locked status of a namespace.

```
[user.hostname]> namespace
[user.hostname] namespace> list
NAME
[user.hostname]-1
[user.hostname]-3
[user.hostname] namespace> select [user.hostname]-3
[user.hostname] namespace '[user.hostname]-3'> get
  type: Namespace
  name: [user.hostname]-3
  description: (unset)
  failedOver: false
  locked: true <----- LOCKED
  namespaceType: REPLICATION
  reference: NAMESPACE-4
  secureNamespace: false
  tag: 5570be25-dbcf-48c3-b2d2-dd2c65eb98b7
[user.hostname] namespace '[user.hostname]-3'>
```

Verify that the namespace cannot be deleted or failed over.

```
[user.hostname] namespace '[user.hostname]-1'> delete
[user.hostname] namespace '[user.hostname]-1' delete *> commit
  Error: Namespace "[user.hostname]-1" is locked and cannot be deleted.
  Action: Cannot delete a locked namespace.
[user.hostname] namespace '[user.hostname]-1' delete *> discard
[user.hostname] namespace '[user.hostname]-1'> failover
[user.hostname] namespace '[user.hostname]-1' failover *> commit
  Error: Namespace "[user.hostname]-1" is locked and cannot be failed over.
  Action: Cannot failover a locked namespace.
[user.hostname] namespace '[user.hostname]-1' failover *> discard
```

Verify that the replication profile cannot be deleted or modified. Objects can still be added to the replication profile.

```
[user.hostname]> replication spec
[user.hostname] replication spec> select locked-spec-1
[user.hostname] replication spec 'locked-spec-1'> delete
[user.hostname] replication spec 'locked-spec-1' delete *> commit
  Error: The replication profile is locked and cannot be deleted.
  Action: Select an unlocked profile to delete.
[user.hostname] replication spec 'locked-spec-1' delete *> discard
[user.hostname] replication spec 'locked-spec-1'> update
[user.hostname] replication spec 'locked-spec-1' update *> set automaticReplication=true
[user.hostname] replication spec 'locked-spec-1' update *> commit
  Error: The replication profile is locked and cannot be updated.
  Action: Select an unlocked profile to update.
[user.hostname] replication spec 'locked-spec-1' update *> discard
[user.hostname] replication spec 'locked-spec-1'> update
```

```
[user.hostname] replication spec 'locked-spec-1' update *> set
objectSpecification.objects=Untitled/dbname,Group:Untitled
[user.hostname] replication spec 'locked-spec-1' update *> commit
[user.hostname] replication spec 'locked-spec-1'> update
[user.hostname] replication spec 'locked-spec-1' update *> set
objectSpecification.objects=Untitled/dbname
[user.hostname] replication spec 'locked-spec-1' update *> commit
Error: Objects cannot be removed from a locked replication profile.
Action: Select an unlocked profile to update.
[user.hostname] replication spec 'locked-spec-1' update *> discard
[user.hostname] replication spec 'locked-spec-1'>
```

Create a replica retention policy and apply it to the locked namespace.

```
[user.hostname]> policy
[user.hostname] policy> createAndApply
[user.hostname] policy createAndApply *> set policy.type=ReplicaRetentionPolicy
[user.hostname] policy createAndApply *> set policy.duration=6
[user.hostname] policy createAndApply *> set policy.durationUnit=YEAR
[user.hostname] policy createAndApply *> set target=Namespace:[user.hostname]-1
[user.hostname] policy createAndApply *> set policy.name="Six Years"
[user.hostname] policy createAndApply *> get
type: PolicyCreateAndApplyParameters
policy:
  type: ReplicaRetentionPolicy (*)
  name: Six Years (*)
  customized: false
  duration: 6 (*)
  durationUnit: YEAR (*)
  target: Namespace:[user.hostname]-1 (*)
[user.hostname] policy createAndApply *> commit
`POLICY_REPLICA_RETENTION-30
[user.hostname] policy>
```

Verify that the replica retention policy cannot be deleted or modified.

```
[user.hostname] policy> select POLICY_REPLICA_RETENTION-30
[user.hostname] policy 'Six Years'> delete
[user.hostname] policy 'Six Years' delete *> commit
Error: The replica retention policy "Six Years" could not be removed because the
target namespace "[user.hostname]-1" is locked.
[user.hostname] policy 'Six Years' delete *> discard
[user.hostname] policy 'Six Years'> update
[user.hostname] policy 'Six Years' update *> set duration=4
[user.hostname] policy 'Six Years' update *> commit
Error: The replica retention policy "Six Years" could not be modified because the
target namespace "[user.hostname]-1" is locked.
[user.hostname] policy 'Six Years' update *> discard
[user.hostname] policy 'Six Years'> unapply
```

```
[user.hostname] policy 'Six Years' unapply *> set target=Namespace:[user.hostname]-1
[user.hostname] policy 'Six Years' unapply *> commit
    Error: The replica retention policy "Six Years" could not be removed because the
target namespace "[user.hostname]-1" is locked.
[user.hostname] policy 'Six Years' unapply *> discard
[user.hostname] policy 'Six Years'>
```

Backup and Recovery Strategies for the Delphix Engine

As a software virtual appliance, Delphix leverages features of the storage, hypervisor, and appliance infrastructure to provide for recovery in the event of failure. These topics walk through the process of evaluating requirements and defining a solution. This process depends on the requirements and features of the environment in which the Delphix Engine is deployed.

- [Backup and Recovery Requirements](#)
- [Backup Solution Implementation](#)
- [Deployment Architecture](#)
- [Mapping Requirements to Solutions](#)

Backup and Recovery Requirements

This topic describes determining requirements for infrastructure failure modes and recovery.

Before devising a strategy, you must first have a set of requirements to evaluate possible solutions. What failures are you trying to protect against, and what are your recovery goals in the event of failure?

Failure Points

Before devising a strategy, you must first have a set of requirements by which the resulting solution can be evaluated. What failures are you trying to protect against, and what are your recovery goals in the event of failure?

Physical Server Failure

The Delphix Engine runs within the VMware ESX hypervisor, which itself is running on a physical machine. Failure of that physical machine will affect the Delphix Engine, as well as any other virtual machines running on that server. The failure is isolated to that particular server and is not the result of a larger, site-wide failure.

- **Recommendation:** ESX Clustering

Storage Failure

The Delphix Engine uses LUNs from a storage array provided through the VMware hypervisor. The storage array may have redundant disks and/or controllers to protect against single points of failure within the array. However, the Delphix Engine can still be affected by a failure of the entire array, the SAN path between the Delphix Engine and the array, or by a failure of the LUNs in the array that are assigned to the Delphix Engine.

- **Recommendation:** Replication

Site Failure

When an entire site or datacenter goes down, all servers, storage, and infrastructure are lost. This will affect not only the Delphix Engine but any production databases and target servers in the datacenter.

- **Recommendation:** Replication

Administrative Error

If an administrator mistakenly deletes a VDB or takes some other irreversible action, there is no method of recovery built into the Delphix Engine.

- **Recommendation:** Snapshots

Recovery Objectives

Once infrastructure fails, some amount of work is required to restore the Delphix Engine to an operational state. Clients won't have access to the Delphix Engine during this time, and the point to which the system is recovered is dependent on the mechanism being used. These qualitative aspects of recovery can be captured by the following metrics. As these metrics are often directly associated with cost, it is important to think not just about the desired metrics, but also the minimum viable goals.

Recovery Point Objective (RPO)

The RPO is the acceptable amount of data that can be lost in the event of a failure. For example, if backups are taken once a day, then at most 24 hours of data will be lost if the system fails immediately before a regularly scheduled backup.

Recovery Time Objective (RTO)

The RTO is the time required to restore the system to an operational state after a failure. For example, recovery may require restoring data from a backup, followed by some number of manual steps to recreate the configuration in the new system. RTO is equivalent to the downtime experienced by clients.

Recovery Time Granularity (RTG)

The granularity of the recovery time is the specificity by which you can select a particular point in time from the past to restore the system. For example, VM snapshots taken every hour provide no way to restore to a point in time between those snapshots.

Related Topics

- [Backup and Recovery Strategies for the Delphix Engine](#)

Backup Solution Implementation

This topic describes the tradeoffs involved with backup and recovery solutions.

With the exception of clustering, solutions can be implemented using features at both the storage and hypervisor layer. Choosing the right technology requires understanding both your requirements and what infrastructure is in use in your environment. The following sections outline some basic choices and the tradeoffs involved.

Clustering

VMware [vSphere's high](#) availability provides the ability to have a VM configuration shared between multiple physical ESX servers. Once the storage has been configured on all physical servers, any server can run the Delphix Engine VM. This allows ESX clusters to survive physical server failure. In the event of failure, the VM is started on a different server and appears to clients as an unexpected reboot with non-zero but minimal downtime. Depending on the length of the outage, this may cause a short pause in I/O and database activity, but longer outages can trigger timeouts at the protocol and database layers that result in I/O and query errors. Such long outages are unlikely to occur in a properly configured environment.

Automatic detection of a failure in an HA environment does not work in all circumstances, and there are cases where the host, storage, or network can hang such that clients are deprived of access, but the systems continue to appear functional. In these cases, a manual failover of the systems may be required.

When configuring a cluster, it is important to provide standby infrastructure with equivalent resources and performance characteristics. Asymmetric performance capabilities can lead to poor performance in the event of a failover. In the worst case of an over-provisioned server, this can cause widespread workload failure and the inability to meet performance SLAs.

Snapshots

VMware provides [storage-agnostic snapshots](#) that are managed through the [VMware Snapshot Manager](#). The use of VMware snapshots can, however, cause debilitating performance problems for write-heavy workloads due to the need to manage snapshot redo-log metadata. In order to provide an alternative snapshot implementation, while retaining the existing management infrastructure, VMware has created an API to allow storage vendors to supply their own snapshot implementation. This is only supported in ESX 5.1. Furthermore, the array must support the [vStorage APIs](#). Consult the [VMware documentation](#) for supported storage solutions and the performance and management implications.

Storage-based snapshots, by virtue of being implemented natively in the storage array, typically do not suffer from such performance problems and are preferred over VMware snapshots when available. When managing storage-based snapshots, it is critical that all LUNs backing a single VM be part of the same consistency group. Consistency groups provide write order consistency across multiple LUNs and allow snapshots to be taken at the same point in time across the LUNs. This must include all VM configuration, system VMDKs, and VMDKs that hold the dSources and VDBs. Each storage vendor presents consistency groups in a different fashion; consult your storage vendor documentation for information on how to configure and manage snapshots across multiple LUNs.

In the event of a snapshot recovery becoming required, ensure that the Delphix Engine VM is powered off for the duration of the snapshot recovery. Failure to do so can lead to filesystem corruption as you're changing blocks underneath a running system.

Replication

[Site Recovery Manager](#) (SRM) is a VMware product that provides replication and failover of virtual machines within a vSphere environment. It is primarily an orchestration framework, with the actual data replication performed by a native VMware implementation, or by the storage array through a storage replication adapter (SRA). A list of

supported SRAs can be found in the [VMware documentation](#). There is some performance overhead in the native solution, but not of the same magnitude as the VMware snapshot impact. SRAs provide better performance but require that the same storage vendor be used as both source and target, and require resynchronization when migrating between storage vendors.

Storage-based replication can also be used in the absence of SRM, though this will require manual coordination when re-configuring and starting up VMs after failover. The VM configuration, as well as the storage configuration within ESX, will have to be recreated using the replicated storage.

The Delphix Engine also provides native replication within Delphix. This has the following benefits:

- The target system is online and active
- VDBs can be provisioned on the target from replicated objects
- A subset of objects can be replicated
- On failover, the objects are started in a disabled state. This allows configuration to be adjusted to reflect the target environment prior to triggering policy-driven actions.
- Multiple sources can be replicated to a single target

Note that the Delphix Engine currently only replicates data objects (dSources and VDBs) and environments (source and target services). It does not replicate system configuration, such as users and policies. This provides more flexibility when mapping between disparate environments, but requires additional work when instantiating an identical copy of a system after failover.

Backup

There is a large ecosystem of storage and VM-based backup tools, each with its own particular advantages and limitations. VMware provides [Data Protector](#), but there are [size limitations](#) (linked to a maximum of 2TB of deduped data) that make it impractical for most Delphix Engine deployments. Most third-party backup products, such as Symantec NetBackup, EMC Networker, and IBM Tivoli Storage Manager, have solutions designed specifically for the backup of virtual machines. Because the Delphix Engine is packaged as an appliance, it is not possible to install third-party backup agents. However, any existing solution that can back up virtual machines without the need for an agent on the system should be applicable to Delphix as well. Check with your preferred backup vendor to understand what capabilities exist.

Some storage vendors also provide a native backup of LUNs. Backup at the storage layer reduces overhead by avoiding data movement across the network but loses some flexibility by not operating within the VMware infrastructure. For example, recreating the VM storage configuration from restored LUNs is a manual process when using storage-based recovery.

Related Topics

- [Backup and Recovery Strategies for the Delphix Engine](#)


Deployment Architecture

This topic describes components of the Delphix deployment architecture.

Delphix operates in a virtual environment with several core systems working in concert, each with its own set of capabilities. Understanding this architecture is critical in evaluating how solutions can be applied across the components, and the tradeoffs involved.

Architectural Components

This diagram illustrates Delphix's recommended best practices for deploying the Delphix Engine in a VMware environment:



This architecture is designed to isolate I/O traffic to individual LUNs while using the most commonly deployed VMware components. In this example, each VMDK file is placed in a separate VMFS volume. Each volume is exported to every node in the ESX cluster, allowing the Delphix Engine to run on any physical host in the cluster.

Fault Recovery Features

Across the recommended deployment architecture there are three key components in play: Delphix Engine, VMware, and storage. Each of these provides different failure handling capabilities, which can be roughly grouped into the following areas.

Server Clustering

Clustering provides a standby server that can take over in the event of failure. A given clustering solution may or may not provide high availability guarantees, though all provide failover capabilities, provided that an identical passive system is available.

Snapshots

Snapshots preserve a point-in-time copy of data that can be used later for rollback or to create writable copies. Creating a snapshot is typically low cost in terms of space and time. Because they use the storage allocated to the array, snapshots restore quickly, but they do not protect against failures of the array.

Replication

Data replication works by sending a series of updates from one system to another in order to recreate the same data remotely. This stream can be synchronous, but due to performance considerations is typically asynchronous, where some data loss is acceptable. Replication has many of the same benefits of backup, in that the data is transferred to a different fault domain, but has superior recovery time given that the data is maintained within an online system. The main drawback of replication is that the data is always current - any logical data error in the primary system is also propagated to the remote target. The impact of such a failure is less when replication is combined with snapshots, as is often the case with continuous data protection (CDP) solutions.

Backup

Like snapshots, backup technologies preserve a point-in-time copy of a storage dataset, but then move that copy to offline storage. Depending on the system, both full and incremental backups may be supported, and the backup images may or may not be consistent. Backup has the advantage that the data itself is stored outside the original fault domain, but comes at high cost in terms of complexity, additional infrastructure, and recovery time.

Related Topics

- [Backup and Recovery Strategies for the Delphix Engine](#)

Mapping Requirements to Solutions

This topic describes how to map from backup and recovery requirements to solutions.

With requirements and detailed knowledge of the deployment architecture, we can map to solutions tailored for the features provided by the underlying infrastructure.

Feature Capabilities

Based on these failure points and recovery features, you can use the following table to map requirements to architectural components: VMware (V), Delphix (D), or storage (S). This can drive implementation based on infrastructure capabilities and recovery objectives.

	Fault Recovery Features			
Failure Point	Clustering	Snapshots	Replication	Backup
Server Failure	V	-	V S D	-
Storage Failure	-	-	V S D	V S
Site Failure	-	-	V S D	V S
Administrative Error	-	V S	-	V S

Recovery Point Objective (RPO)

Feature	Time	Description
Clustering	Zero	All changes committed to disk are automatically propagated to the new server. Any pending changes in memory are lost.
Replication	Near zero	Most solutions offer scheduled replication, but many can offer continuous replication with near-zero data loss.
Snapshots	Snapshot period (for example, one hour)	Given their relatively low cost, snapshots tend to be taken at a higher frequency than a traditional backup schedule.

Feature	Time	Description
Backup	Backup period (for example, one day)	Backup policies can be configured in a variety of ways, but even with incremental backups, most deployments operate no more frequently than once a day because of the cost of full backups, and the impact of incremental backups on recovery time.

Recovery Time Objective (RTO)

Feature	Time	Description
Clustering	Near zero	VM clustering with the Delphix Engine provides near-zero downtime in the event of failure, but clients may be briefly paused or interrupted.
Replication	15 minutes	The target side environment is kept in hot standby mode, so it is relatively quick to switch over to the target environment. Depending on the scope of the failure, however, some configuration information may need to be changed on the target side prior to enabling operation.
Snapshots	15 minutes	The Delphix Engine can be rolled back to a previous state. Changes made to systems external to the Delphix Engine (for example, deleting a VDB) can cause inconsistencies after rollback.
Backup	Hours or days	Restoring a full backup can be very time-consuming. In addition to having to read, transfer, and write all of the data, the same process will need to be run for each incremental backup to reach the objective point.

Recovery Time Granularity

Feature	Granularity	Description
Clustering	None	Only the current system state can be recovered.
Replication	None	Only the nearest replicated state can be recovered unless combined with snapshots.
Snapshots	Snapshot period (for example, one hour)	Determined by the snapshot schedule.

Feature	Granularity	Description
Backup	Backup period (for example, one day)	Determined by the backup schedule.

Related Topics

- [Backup and Recovery Strategies for the Delphix Engine](#)

Delphix Replication

This section covers the following topics:

- [Replication Overview](#)
- [Replication Concepts](#)
- [Replication Use Cases](#)
- [Configuring Replication](#)
- [Managing Replicated Objects](#)
- [Controlled Failover](#)
- [Uncontrolled Failover](#)
- [Replicas and Failover](#)
- [Provisioning from Replicated Data Sources or VDBs](#)
- [Replication User Interface](#)

Replication Overview

Delphix allows data objects to be replicated between Delphix Engines. Prior to version 5.3.3, these Delphix Engines had to be running identical Delphix versions, but could otherwise vary in terms of other configurations. With versions 5.3.3 and above, engines can now be on different versions, which is further detailed in the [Forward Compatible Replication \(FCR\)](#) section. In the event of a failure that destroys the source engine, you can bring up the target engine in a state that matches the source. In addition, you can provision VDBs from replicated objects, allowing for the geographical distribution of data and remote provisioning.

You can run replication ad hoc, but it is typically run according to a predefined schedule. After the initial update, each subsequent update sends only the changes incurred since the previous update. Replication does not provide synchronous semantics, which would otherwise guarantee that all data is preserved on the target engine. When there is a failover to a replication target, some data is lost, equivalent to the last time a replication update was sent.

Replication is generally not suited for high-availability configurations where rapid failover (and failback) is a requirement. Failing over a replication target requires a non-trivial amount of time and is a one-way operation; to failback requires replicating all data back to the original source. For cases where high availability is necessary, it is best to leverage features of the underlying hypervisor or storage platform. For more information on how to evaluate the use of Delphix Engine replication for your data recovery requirements, see the topics under [Backup and Recovery Strategies for the Delphix Engine](#).

Forward Compatible Replication (FCR)

Delphix Virtualization supports the ability to replicate to a Delphix Target Engine running on a higher version. To do so, FCR has a few requirements and restrictions to consider:

- FCR is supported for replication jobs starting from a **source engine** running 5.3.0.0 and beyond.
- The **target engine** must be running 5.3.3.0 and beyond.
- As of 6.0.10.0, the FCR replication operation can not go beyond engine versions released more than 12 months apart.
- FCR supports replicating between major versions as long as those specific versions are no more than a year apart.

Examples of supported and not supported FCR configurations:

Supported

- 6.0.3.0 to 6.0.3.0 (same version)
- 6.0.3.1 to 6.0.3.1 (same version)
- 6.0.3.0 to 6.0.3.1 (higher patch version)
- 6.0.3.1 to 6.0.5.0
- 6.0.3.0 to 6.0.9.0
- 6.0.4.0 to 6.0.10.0

Not Supported

- 5.2.5.0 to 6.4.0.0 (source version not compatible with FCR)
- 5.3.0.0 to 5.3.2.0 (target version not compatible with FCR)
- 6.0.9.0 to 6.0.8.0 (source version higher than the target)
- 6.0.0.0 to 6.0.10.0 (there is more than a year between release dates for those versions)

Exceptions

Some of the newer versions of 5.3.x are not compatible with the early versions of 6.0.x.0 (6.0.0.0 and 6.0.1.0).

- 5.3.7.0 and 5.3.7.1 are not compatible with 6.0.0.0
- 5.3.8.0 and 5.3.8.1 are not compatible with 6.0.0.0
- 5.3.9.0 is not compatible with 6.0.0.0
- 5.3.9.0 is not compatible with 6.0.1.0 and 6.0.1.1

Replication Matrix

The following table lists all the Delphix Engine versions that a user can replicate to the required or the latest version.

Key:

Color	Supported?
	Yes
	No

Replicate To Target																												
R e p l i c a t e F r o m S o u r c e	5	5	5	5	5	5	5	5	5	5	5	5	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6	6
	3	3	3	3	3	3	3	3	3	3	3	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	1	2	3	4	5	6	7	7	8	8	9	0	1	1	2	2	3	3	4	4	4	5	6	6	7	8	8	9
0	0	0	0	0	0	0	0	1	0	1	0	0	0	1	0	1	0	1	0	1	2	0	0	1	0	0	1	0
5.3.0.0																												
5.3.1.0																												

Replicate To Target

Replication Features

As virtual appliances, it is possible to backup, restore, replicate, and migrate data objects between Delphix Engines using features of VMWare and the underlying storage infrastructure. Data objects include groups, dSources, VDBs, Self-Service (Jet Stream) data templates and data containers, and associated dependencies. In addition to the replication capabilities provided by this infrastructure, native Delphix Engine replication provides further capabilities, such as the ability to replicate a subset of objects, replicate multiple sources to a single target, and provision VDBs from replicated dSources and VDBs without affecting ongoing updates. The topics under [Backup and Recovery Strategies for the Delphix Engine](#) provide more information on how to evaluate features of the Delphix Engine in relation to your backup and recovery requirements.

Replication is configured on the source Delphix Engine and copies a subset of dSources and VDBs to a target Delphix Engine. It then sends incremental updates manually or according to a schedule. For more information on configuring replication, see [Configuring Replication](#).

You can use replicated dSources and VDBs to provision new VDBs on the target side. You can refresh these VDBs to data sent as part of an incremental replication update, as long as you do not destroy the parent object on the replication source. For more information, see [Provisioning from Replicated Data Sources or VDBs](#).

During replication, replicated dSources and VDBs are maintained in an alternate replica and are not active on the target side. In the event of a disaster, a failover operation can break the replication relationship. For more information on how to activate replicated objects, see [Replicas and Failover](#).

Replication Details

When you select objects for replication, the engine will automatically include any dependencies, including parent objects, such as groups, and data dependencies such as VDB sources. This means that replicating a VDB will automatically include its group, the parent dSource, and the group of the dSource, as well as any environments associated with those databases. When replicating an entire engine, all environments will be included. When replicating a database or group, only those environments with the replicated databases are included.

During replication, the Delphix Engine will negotiate an SSL connection with its server peer to use TLS_RSA_WITH_AES_256_CBC_SHA256 as the cipher suite and TLSv1.2 as the protocol.

Only database objects and their dependencies, as well as certain non-database objects, are copied as part of a backup or replication operation, including:

- dSources
- VDBs
- Groups
- Self-service (Jet Stream) Data Templates and Data Containers
- Environments
- Environment configuration (users, database instances, and installations)
- Delphix users, roles, permissions, and authorizations
- Policies
- Database configuration templates

The following objects are NOT copied as part of a backup or replication operation:

- Events and faults
- Job history
- System services settings, such as SMTP
- Hook templates
- Alert profiles

After failover, you must recreate these settings on the target.

On-Premises Replication To Azure/OCI/GCP/Hyper-V

Replicating from on-premises engines with an underlying storage block size of 512B will experience disk usage inflation when replicating to target engines with different underlying block sizes. Azure, GCP, Hyper-V, and OCI are known to have 4K block sizes and therefore will require extra disk capacity when receiving replication from an on-premises engine. This behavior is due to the fact that the underlying storage block size is different (512B vs. 4K) between the two Delphix Engines (one on Prem, one on Azure/OCI/GCP/Hyper-V), resulting in a lower compression rate on the replication target. It is expected that 1.5-1.6x the amount of space is taken from objects on-premises in these cases.

Resumable Replication

Resumable replication enhances the current replication feature by allowing you to restart large, time-consuming initial replication or incremental updates from an intermediate point. A single replication instance can fail for a number of environmental and internal reasons. Previously, when you restarted a failed replication instance, replication required a full resend of all data transmitted prior to the failure. With resumable replication, no data is retransmitted. Replication is resumable across machine reboot, stack restart, and network partitions. The resumable replication feature is fully automated and does not require or allow any user intervention.

For example, suppose a replication profile has already been configured from a source to a target. A large, full send begins between the two that is expected to take weeks to complete. Halfway through, a power outage at the data center that houses the source causes the source machine to go down and only come back up after a few hours. On startup, the source will detect a replication was ongoing, automatically re-contact the target, and resume the replication where it left off. In the user interface (UI) on the source, the same replication sends job will appear as active and continue to update its progress. However, in the UI of the target, a new replication receives job will appear but will track its progress as a percentage of the entire replication.

In 4.1 and earlier releases, the replication component would always clean up after failed jobs to ensure that the Delphix Engine was kept in a consistent state and that no storage was wasted on unused data. With the addition of reusability, the target and source can choose to retain a partial replication state following a failure to allow future replications to complete from that intermediate point. In the current release, the target and source will only choose to retain partial replication state following failures that leave them out of network contact with each other – for example, source restart, target restart, or network partition. Once network contact is re-established, the ongoing replication will be automatically detected and resumed.

Replication will not resume after failures that leave the source and target connected. For example, if a storage failure on the target, such as out-of-space errors, causes a replication to fail, the source and target remain connected. As a result, the Engine will discard state data associated with the failed Replication operation. Nonetheless, resumable replication would begin during a source reboot, a target reboot, and a network partition.

Replicating Delphix Self-Service Templates

Templates can now be replicated and accessed on the target engine via Delphix Self-Service (Jet Stream). Replicated templates can be replicated into the target space with or without their containers. On the new target engine, the newly created replicated template can be used to create new containers that are assigned to users. You cannot change the replicated template's name or the names of the containers from which it was replicated.

 Any containers that were replicated over with the template cannot be used to start, stop, etc until they are disconnected to their parent containers in the source engine during the failover operation.

Replication of Non-data Objects

As of 6.0.5.0, the replication of non-data objects is supported.

Non-data objects refer to:

- Delphix users, roles, permissions, and authorizations
- Policies
- Database configuration templates

These objects will not be presented as selectable objects when creating a replication spec, but will instead be passively included by association, the same way environments are. Replication of non-data objects will follow these rules:

- If the entire engine is replicated, all non-data objects will be included
- If specific data objects are replicated, all associated non-data objects will be included

For example, when a container is replicated, then policies that apply to that container, as well as users who have authorizations on that container, will all be included for replication.

Replicated users will be shown on the Received Replicas page on the Delphix Engine UI.

Also as of 6.0.5.0, the [automatic conflict resolution](#) option is chosen by default; it is the recommended way of handling the failover of non-data objects, as non-data objects are expected to cause collisions. Automatic conflict resolution will resolve non-data object conflicts by the following rules:

- Users will be consolidated if both username and email match. Otherwise, the replicated user will be renamed.
- Roles, permissions, and authorizations will be consolidated
- Policies and database configuration templates will be renamed, except for "None" type policies, which will be consolidated. Following failover, replicated policies will continue to apply to the same replicated data objects.

Related Topics

- [Replication Concepts](#)
- [Replication Use Cases](#)

Replication Concepts

Delphix replication allows you to copy objects from one Engine (referred to as a source Engine) to another Engine (referred to as a target Engine).

Replication recreates objects from the source Engine onto the target Engine in a **replica**, also known as a **namespace**, that preserves object relationships and naming on the target Engine without interfering with its active objects. Objects within a replica are read-only and disabled until the replica is failed over, at which point they can be activated. VDBs and dSources within a replica can be used as the source for provisioning new VDBs. Below are key concepts for replication that will be explained in detail:

- **Received Replicas or Namespaces:** Once replication is complete, the target Engine will create a received replica, also known as a namespace. This is a copy of objects that are related via a grouping.
- **Failover and Conflict Resolution:** Certain objects may require changes to resolve conflicts prior to completing a replication failover. Names of replicated objects should be unique for replication to complete successfully
- **Enabling Databases and Environments:** Once replication takes place, you must enable the databases and environments on the target engine. This step ensures the configuration of these objects is correct and will ensure that replication has correctly copied the necessary object relationships from the source engine.

Received Replicas or Namespaces

A replica contains a set of replicated objects, such as dSources and VDBs. These objects are read-only and disabled while replication is ongoing. You may view replicated objects both in the Delphix Engine UI as well as the CLI. To view received replicas:

1. On the target Engine, navigate to 'System'. Then select Replication.
2. Under Received Replicas, select the replica.

On this screen, you can browse the contents of replicas, as well as failover or delete individual replicas.

Deleting or failing over a replica will break any link with the replication source. Subsequent incremental updates will fail, requiring the source to re-establish replication. Failover should only be triggered when no further updates from the source are possible, as in a disaster scenario. Various replication use cases are also described in [Replication Use Cases](#).

Multiple replicas can exist on the system at the same time. Active objects can exist in the system alongside replicas without interfering with replication updates. You can also use VDBs and dSources within a replica as a source when provisioning. For more information, see [Provisioning from Replicated Data Sources or VDBs](#).

Failover and Conflict Resolution

To activate the objects in a replica, you must first fail over the replica. This will disconnect the replication connection and move the objects to the live system, where they can be actively used.

During the replication failover, there may be objects that conflict between the source Engine and the target Engine. For example, 'groups' will conflict if the failing over group has the same name as a group in the live system, as well as environments, dSources, and VDBs. By default, active objects with conflicting names will cause an error at the time of failover. In these scenarios, you must rename the active objects before the failover operation can complete successfully.

Given that conflicting names prevent failover from succeeding, best practices in a disaster recovery situation are to leave the target system completely passive with no active objects until the time that failover is required.

Once a replica is failed over, the objects are active but will be automatically disabled. The next section describes enabling these objects after replication is complete.

Enabling Databases and Environments

Objects may refer to states (IP addresses, mount paths, etc.) that differ between the source and target Engines. Because of this, all databases and objects within a replica are automatically disabled after a failover. This allows the administrator to alter configuration prior to enabling databases and environments, without the system inadvertently connecting to invalid systems.

After failover is complete, you must explicitly enable all dSources, VDBs, and environments. If you need to change any configuration for the target environment, you can do so prior to enabling the objects. In the event that a failing-over environment is consolidated with a live system environment, it must be refreshed before all of its databases can be used.

Replication Use Cases

Replication allows you to move Delphix objects such as dSources and virtual databases (VDBs) between Delphix Engines. These topics describe how you can use replication to meet different use cases:

- **Replicating to the Public Cloud:** With Delphix replication you can send data from engines deployed on-premise to the public cloud. In this case, you may have a Delphix engine in the production zone to ingest data securely with replication set up to the public cloud for development access and rapid, on-demand testing.
- **Disaster Recovery:** In the event of a disaster, you may need to failover your engine. Delphix replication enables you to configure a failover Delphix Engine to preserve the data and Delphix objects from the source engine for disaster recovery.
- **Geographically Distributed Development:** Often, development teams access data from all over the world. In this scenario, you may want to replicate data to be local to the developers who require access. With Delphix replication, you can provide developers with local access to Delphix datasets.
- **Data Migration:** Delphix supports simple migration of data and resources between Delphix Engines. There may be cases when you need to consolidate workloads between different Delphix Engines. With replication, you can easily migrate data as needed.

With Delphix replication, you can achieve the ideal Delphix deployment:

1. Ingesting data in production with a Delphix Engine deployed within the production zone.
2. Masking the data in production using [Delphix Masking](#).
3. Replicating the masked data to a non-production environment, using Selective Data Distribution (SDD). SDD enables you to securely replicate masked data without compromising sensitive data from the source engine. For more information, read [Selective Data Distribution Overview](#).

For more information on a few examples of Replication use cases, view the sections below.

Replicating to the Public Cloud

Enterprises will usually have the separate infrastructure for production systems and non-production development environments. For example, in the hybrid-cloud model, on-premise data centers are used to allow the company maximum ownership of these systems, while the public cloud is leveraged by development and automated testing teams to accelerate software development.

Delphix Engine Deployed in the Public Cloud

- Once a Delphix Engine has been deployed in the public cloud (using a supported cloud platform as described in [Deployment](#)) you can begin replicating from any source engine to that target engine.
- This enables you to provide access to production data from a Delphix Engine deployed on-premise to a Delphix Engine deployed in the cloud platform of your choice.

Disaster Recovery

Replication is often used to provide recovery in the event of a disaster, where a data center or site becomes completely unavailable. To prepare for this, you may configure a failover Delphix Engine, which we will refer to as the *replication target Engine*. This target Engine will regularly receive replication updates from the original, or source, Engine so that if the source ever becomes unavailable, the target can be activated immediately.

Configuration Steps

- **Passive Target Engine**

- For disaster recovery, the target engine should be kept in a passive state until the source engine is lost. This prevents failover conflicts that may occur during the replication process.
- At this point, a failover is performed that breaks subsequent replication updates and activates objects so that you can manage them on the target side.
- **Same Configuration for Source and Target Engines**
 - Target hosts and systems should exist at the target that matches the configuration of those at the primary engine. For example, if the source engine has two Red Hat environments discovered with four Oracle databases, the target engine should have exactly two Red Hat environments and four Oracle databases as well.
 - The failover engine should be provisioned with identical resources as the primary engine. For example, both engines should have the same number and types of disks attached as storage.
 - Finally, both engines should have the same network and storage topologies.

Failover Object Management

Once a failover occurs, there are two scenarios that will affect how you manage replicated objects, which include dSources, VDBs, and Environments.

- **Failure of Infrastructure Running the Delphix Engine Only**
 - You can enable dSources and VDBs and reconnect to the original environments that existed on the Source Engine.
 - You can reconfigure environments on the target Engine prior to failover as well.
- **Failure of Infrastructure Running the Delphix Engine and Delphix-connected Environments**
 - Environments will then have to be adjusted to point to the new systems on the target side.
 - If there is not a 1:1 mapping, then you can migrate the VDBs to new systems on the target, and you can detach dSources and attach them to the standby system in the target environment.

Follow the best practices below to simplify failover and meet performance expectations in the event of a disaster:

- To the extent possible, the failover Engine should mirror the primary Engine
- Maintain a 1:1 relationship between source and targets. All data-related objects such as dSources, VDBs, and Environments as well as their configurations such as users and policies.
- The target Engine should remain passive and not be actively used for other workloads

Geographically Distributed Development

Organizations often have development teams distributed across different networks as well as different geographical locations. To improve performance or even meet security requirements, it may be necessary to replicate data from one location to another. The Delphix Engine allows for VDBs to be provisioned from replicated dSources and VDBs, as described in [Provisioning from Replicated Data Sources or VDBs](#).


This use case differs from **Disaster Recovery** as replication is never broken and failover is never performed. You can refresh remote VDBs as long as the parent objects continue to exist on the source. If they are deleted, then remote VDBs will continue to exist but cannot be refreshed or updated from their original source VDB.

Configuration Steps

- **Avoid Heavy VDB Workloads on the Source:** Because each replication stream induces load on the source system:
- Minimize the number of simultaneous replication updates. Each source engine can support replicating to multiple target engines, but you should try to keep simultaneous updates to under three target engines per source.
- If possible, avoid heavy VDB workloads on the Source Engine

- **The permanence of Source Objects:** Provision only from sources that are effectively permanent, since replicated VDBs cannot be refreshed once the source is deleted. If you delete a source object you will need to re-replicate the VDBs if they need to be refreshed.
- **Additional Storage Capacity on the Target Engine:** Provision additional storage capacity on the target Engine
- Remotely provisioned VDBs can consume shared storage (via NFS mounts) on the target even when the parent is deleted on the source


This use case supports more complex models such as 1-to-many and many-to-1. However, replication can only replicate objects that exist in the primary namespace. Consider a Delphix deployment with three Engines: Engine A, Engine B, and Engine C. The following workflow **is supported**:



1. **Engine A Replicates to Engine B**
2. **Provision a Replica VDB on Engine B**
3. **Engine B Replicates to Engine C**

In this scenario, you can replicate from Engine A to Engine B. Then, on Engine B provision a replica VDB into the primary namespace. Finally, add that object into a replication specification to replicate to Engine C. **Note:** Engine B can only replicate objects that exist in the primary (live) namespace to Engine C.

It is important to note the interim step of provisioning an on Engine B is required. For example, the following usage **is not allowed**:



1. **Engine A Replicates to Engine B**
2. **Engine B Replicates to Engine C**

You can replicate from Engine A to Engine B. But on Engine B, it is not possible, nor supported, to create a replication specification with objects in the replication namespace that are desired to be replicated to Engine C.

Data Migration

You can use replication to perform a one-time migration of resources, such as virtual databases or environments, from one Delphix Engine to another. While the hypervisor provides tools to move virtual appliances between physical systems, there are times when migration is necessary, such as:

- Migrating between different physical storage
- Consolidating or distributing workloads across Delphix Engines

In these cases, you can use replication to copy a subset of objects across different topologies.

For migration, follow these best practices:

- Send full updates, followed by incremental updates, until the time required for incremental updates meets your downtime window
- Disable all objects to be migrated on the source, to ensure that they are not actively changing
- Send a final incremental update before failing over the target Engine
- After failover, delete any migrated objects on the source, or the entire Engine

Configuring Replication

This topic describes how to configure data replication between Delphix Engines. Replication is configured with *Replication Profiles* that contain options such as the replication schedule, the hostname of the target Engine, and the selected objects that will be replicated.

Requirements

- **Version Requirements:** the replication target can be on the same or newer version than the replication source.
- **Engine Communication:** The target Delphix Engine must be reachable from the source Engine.
- **Storage Allocation:** The target Delphix Engine must have sufficient free storage to receive the replicated data.
- **Privilege Requirements:** The user in the replication profile must have administrative privileges on the source and the target engines.

Configuring the Network

Delphix Replication uses a private network protocol to communicate between two Delphix Engines. You may specify a network interface to run replication by configuring routing to direct traffic over a particular interface.

The replication network protocol uses TCP port 8415. If there is a firewall between the source and target that is blocking this port, then there are two possible solutions:

1. Enable port 8415 on the firewall in order to allow connections to this port from the source to the target.
2. Replication can connect through a SOCKS proxy if one exists. Configure the SOCKS proxy address and port by connecting to the command-line interface (CLI) as a system administrator user and navigating to "service proxy" to update the SOCKS configuration.

Port 1080

SOCKS port 1080 is used by default but can be overridden

Replication can connect through a SOCKS proxy if one exists. Configure the SOCKS proxy address and port by connecting to the command-line interface (CLI) as a system administrator user and navigating to "service proxy" to update the SOCKS configuration. Example:

Example of a SOCKS Proxy

```
d\px-engine> service proxy
d\px-engine service proxy> update
d\px-engine service proxy update *> set socks.enabled=true
d\px-engine service proxy update *> set socks.host=10.2.3.4
d\px-engine service proxy update *> set socks.username=someuser
```

```
dlpx-engine service proxy update *> set socks.password=somepassword
dlpx-engine service proxy update *> commit
dlpx-engine service proxy> get
  type: ProxyService
  https:
    type: ProxyConfiguration
    enabled: false
    host: (unset)
    password: (unset)
    port: 8080
    username: (unset)
  socks:
    type: ProxyConfiguration
    enabled: true
    host: 10.2.3.4
    password: *****
    port: 1080
    username: someuser
```

Configuring the Replication Source Delphix Engine

1. On the source Delphix Engine, click **System**, then **Replication**.
2. In the left-hand navigation section, click **Create Profile**.
3. Enter the following required fields:
 1. Name of the replication profile
 2. The hostname or IP address for the target Delphix Engine.

Replication Profile Options

There are several configuration options for your replication profiles. These give you more granular control on options such as when replication will run, how much bandwidth it may use, and which objects are replicated. Details for each option are described below.

 The following configurable options are static and can not be configured at run-time. You can set these configurations while a replication spec is being executed, but the values will be applied only after the next execution.

- **Automatic Replication:** A policy to automatically run replication. With this option, you can set up replication based on the schedule you need. Automatic replication allows you to define a policy to automatically run replication. By default, automatic replication is disabled, meaning that you must trigger replication updates manually. To enable automatic replication, click the Enabled checkbox. With this setting, you can enter the frequency and time for replication updates to the target Delphix Engine. Automatic replication uses Quartz, a job scheduling tool (<http://www.quartz-scheduler.org/>), for scheduling, which can be configured via the Advanced option.
- **Traffic Options:** Various traffic and bandwidth options are available. For example, you may want to enable encrypted traffic or limited bandwidth during replication updates.

- **Encrypting Traffic:** By default, replication streams are unencrypted, which provides maximum performance on a secure network. However, this setting allows you to encrypt traffic during replication.

Encrypting Replication

Encrypting the replication stream will consume additional CPU resources and may limit the maximum bandwidth that can be achieved.

- **Network Connections:** Allows setting the number of underlying network connections that can be used by replication.
- **Limiting Bandwidth:** By default, replication will run at the maximum speed permitted by the underlying infrastructure. In some cases, particularly when a shared network is being used, replication can increase resource contention and may impact the performance of other operations. This option allows administrators to specify the maximum bandwidth that replication can consume.
- **Objects Being Replicated:** Select the objects you wish to be replicated from the source engine to the target engine. In the right-hand column, under Objects Being Replicated, you can select the objects you want to replicate. Some selected objects may have dependencies – other objects that will be pulled into replication because they share data.



- This is not guaranteed to be the full set of dependent objects. The full set of objects and their dependents will be calculated at the time of replication.
- You can not configure a Replication Profile without selecting an object and the last object from a Replication Profile can not be removed. If you need to remove the last object from the Replication Profile, you must delete that Replication Profile.
- The object is removed from the replication target namespace only after a subsequent replication job is executed for the associated replication specification.
 - Remove the object from the specification on the source
 - Execute the replication specification
 - The object is removed as a part of the replication job.

The object is not removed while modifying the replication specification. To add back the object, the entire object needs to be sent again.

1. When replicating a group, all dSources and VDBs currently in the group, or added to the group at a later time, will be included.
2. If you select a Delphix Self-Service data template, all data containers created from that template will be included. Likewise, if you select a data container, its parent data template will be included.

3. Regardless of whether you select a VDB individually or as part of a group, the parent dSource or VDB (and any parents in its lineage) are automatically included.
 1. This is required because VDBs share data with their parent object.
 2. In addition, any environments containing database instances used as part of a replicated dSource or VDB are included as well.
 - When replicating individual VDBs, only those database instances and repositories required to represent the replicated VDBs are included. Other database instances that may be part of the environment, such as those for other VDBs, are not included.
 - Non-data objects (Delphix users, roles, permissions, authorizations, policies, database configuration templates) that are associated with selected objects will be automatically included during replication. They will not be shown on the selection interface.

Configuring the Target Delphix Engine

Additional configuration on the target engine is not needed. Replicated objects will appear in an alternate received replica (or namespace) that mirrors the original object layout.

To view replicated objects from the Delphix Engine:

1. Click System, then select Replication.
2. Look under Received Replicas. All replicated objects are read-only until the replica is failed over. For more information about managing replicas and how to activate a replica, see the topics [Replicas and Failover](#) and [Controlled Failover](#).

You can create and manage objects on the target server without affecting subsequent updates, though this can cause conflicts on failover that require additional time to resolve. For disaster recovery use cases, it is recommended to keep the target passive and not create any local objects. This will avoid conflicts and guarantee a smooth failover operation.

Multiple sources can replicate to the same target, allowing for the flexible geographical distribution of data. This is not a recommended practice for disaster recovery, because it increases the probability of conflicts on failover and may oversubscribe resources on the target if multiple replicas are failed over and there is insufficient infrastructure to support the combined workloads.

Related Topics

- [Delphix Replication](#)

Managing Replicated Objects

Once the failover process has completed the last step is to enable the objects on the Delphix Engine that received a replica. There are two major concepts:

1. **SSH Keys:** Public keys need to be re-published on the source, staging, and target environments for the replicated Delphix Engine, since the replication process does not automatically copy SSH keys.
2. **Enabling Replicated Objects:** Objects will need to be re-enabled once replication is complete. This ensures that replication has completed successfully and the configuration of each object is correct.

SSH Keys

If you were previously using a password-less SSH key exchange on your source and/or target environments, you may need to publish the now-primary engine's public key for continued functionality.

1. Click **Manage**.
2. Select **Environments**.
3. Select **any environment**.
4. Under **Environment Users**, click the **edit** icon () next to the defined user.
5. Click **View Public Key** to display the public key for the Delphix Engine.
6. Highlight the public key string (starts with "ssh-rsa") and copy the key to your clipboard (Ctrl+C in Windows).
7. On each source and target host within your defined environments, paste the public key into the environment user's `authorized_keys` file, which is normally located in the user's `~/.ssh/` directory.

Once the public key has been copied to the hosts making up your environments, you are ready to enable the remaining objects.

Enabling Environments

In order to begin using the environments and related dSources and VDBs, the environments must first be started.

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Environments**.
4. Select an **environment** that you want to enable.
5. From the Actions menu (...) select **Enable**. This will initiate jobs that will refresh and enable the environment.
6. Repeat these steps for each environment that you want to enable.

Enabling dSources

In order to again have actively syncing dSources, you must enable them.

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **dSource** you want to enable.
5. From the Actions menu (...) select **Enable**.
6. Click **Enable** to confirm.

Enabling VDBs

The final step in the failover would be to enable any needed VDBs.

1. Login to the Delphix Management application.
2. Click **Manage**.
3. Select **Datasets**.
4. Click the **VDB** you want to enable.
5. Click the **Configuration** tab.
6. From the Actions menu (...) select **Enable**.
7. Click **Enable** to confirm that you want to enable the VDB.

At this time the failover is complete. All objects previously running on the previous source Delphix Engine should now be running off of the target Delphix Engine on which the objects were failed over.

Related Topics

- [Delphix Replication](#)

Controlled Failover

This topic describes the process of failing over a replica. Objects stored in a replica are read-only, and failing over a replica moves the replicated objects to the live system. After a failover, all of the objects will appear in the system as if they had been created locally.

Prerequisites


- A Delphix system that contains a replica is required.
- For an overview of what replicas are and what failover implies, see [Replicas and Failover](#).
- For more information on configuring replication, refer to the [Configuring Replication](#) topics.

Procedure

1. Locate the replica to failover.
 - a. Click **System**.
 - b. Select **Replication**.
In the **Received Replicas** section, you will see the list of replicas. Each replica has a default name that is the hostname of the source that sent the update. If you wish, you can customize these names. Each replica will list the databases and environments it contains.

⚠ If this replica is the result of a replication update, check to see whether or not the source Delphix appliance is still active. If it is still active, then disable any dSource or VDB that is part of the replica being failed over to ensure that only one instance is enabled. You can disable dSources and VDBs by selecting **Datasets**, finding the appropriate database, select **Disable**. After disabling the objects, navigate to **System > Replication** and click **Replicate Now** to get the most recent data to the target environment. If any scripts are accessing the engine, then you must repoint those scripts to the replication target environment to avoid VDBs from being started up in the old replication source environment.

2. Click **Failover** and a confirmation dialog will appear (see screenshot "failover-confirm"). To accept the default failover behavior, confirm the dialog.



This will pause while the replica is failed over.

3. When failover completes, the replica page will update.

4. Apply any configuration changes that are required to customize the objects for the system. This might include updating object states such as IP addresses, mount paths, or credentials. For more details, see the [Replicas and Failover](#) topic.
5. Enable the environments that were failed over.
 - a. Click **Manage**.
 - b. Select **Environments**. The environments that were failed over will be disabled.
 - c. From the Actions menu (...) select **Enable**.
6. Refresh any environments that were consolidated during failover.
 - a. Click the **Refresh** icon for each affected environment.
7. Enable the dSources and VDBs that were failed over.
 - a. Click **Manage**
 - b. Select **Datasets**.
 - c. Select the desired **database**.
8. From the Actions menu (...) select **Enable**.

 If the dSources and VDBs that were failed over belong to a plugin, depending on what version of the plugin existed on the target engine, the plugin may be in an inactive state. The plugin must be moved out of the inactive state before the object can be enabled. For more information see the [Delphix Engine Plugin Management](#) topic.

After Replication Failover

Navigate to **System > Replication** and select the delete button to delete the replication specification.

A failover is similar to a migration operation for TDE-enabled vPDBs. If after the failover the original target hosts are still available, then the vPDB can be enabled on the failed over Delphix Engine without requiring additional manual steps. If the vPDB is to be restored to a different target host, then the same manual steps are required for migration, as the vPDB will now be located on a new target and thus needs the artifact directory, parent Keystore, and merged Keystore available.

Related Topics

- [Replication User Interface](#)
- [Replicas and Failover](#)
- [Configuring Replication](#)
- [Transparent Data Encryption and Delphix](#)

Uncontrolled Failover

 Before starting an uncontrolled failover, keep a note of the source and target engine's IP address and engine UUID. These details are available on the [Delphix Server Setup Application](#)

Uncontrolled Failover Procedure

1. On the source engine, create a replication spec for the dSources and VDBs, by clicking **Replicate Now** to replicate these dSources/VDBs to the target engine.
2. To simulate a crash of the source engine, Power of the source engine.
3. On the target engine, click on **Failover**.
At this point, the target engine will become active and the replicated environments/dSources/VDBs will be seen in the default namespace of the target engine.

The following sections describe the procedures to handle an Uncontrolled Failover for each data platform:

- [Oracle Environment and Data Sources](#)
- [SAP ASE Environment and Data Sources](#)
- [SQL Server Environment and Data Sources](#)

Oracle Environment and Data Sources

After an uncontrolled failover, for Oracle SI and RAC complete the following procedure:

1. Enable dSources and VDBs on the target engine.
dSources will be enabled but the VDBs/vPDBs will fail.
2. To enable VDBs, shutdown the database instance using **SHUTDOWN ABORT**.
3. Remove all the mounts from the target environment with sudo privilege.
 - a. List all the mounts on the target environment using the **mount** command.
 - b. Run **umount -lf <Delphix-mount>** to remove the mounts.
For example: **umount -lf /mnt/provision/VDBOMSR8A1718_FPX/datafile**
4. Remove all the stale oracle processes of the database instance using **kill -p <process id>**.
 - a. Run **ps -ef | grep "DATABASE INSTANCE** to get all the running processes associated with the database instance.
For example **ps -ef | grep VDBOMSR8A1718_FPX**. This command list all the running processes, e.g. **oracle 15236 9227 0 11:01 pts/0 00:00:00 grep --color=auto VDBOMSR8A1718_FPX**
 - b. Run **kill -p <process id>** to kill all the running processes.
For example: **kill -9 15236**
5. Now enabling VDBs and VPDBs will be successful.

Related Topics

- [Uncontrolled Failover](#)

SAP ASE Environment and Data Sources

After an uncontrolled failover has been triggered, on the target/staging host, the SAP ASE dSource/VDBs mounts, corresponding to the source engine, will still be present. This can be checked using the following command on the target host:

```
mount | grep <source-engine-ip>
```

These mounts will be mounted in the user-provided toolkit directory (Example */work*) and these would be of the form **/work/<Engine-UUID>-[staging/vdb]/[datafile/archive/temp]**. The following is an example for a dSource staging database mount **/work/4201763f-2f8d-1c8f-381a-efb180b0328-staging-2/datafile**.

All the mounts will be using the source engine's UUID.

Attempting to '**Enable**' or '**Start**' the ASE dSources/VDBs on the target engine while the NFS mounts are in this stale state, will cause the jobs to hang or timeout or result in a job failure.

To clear off these stale mounts, one of the following options can be used:

Manually unmount

1. Shutdown the concerned SAP ASE instances which host the dSource's staging database and VDBs.
2. List all the dSource/VDB mounts with a **mount | grep <source-engine-ip>**.
3. Remove all the mounts with a **umount -lf <Delphix-mount>**
for example: **umount -lf /work/4201763f-2f8d-1c8f-381a-efb180b0328-staging-1/archive**, this might require sudo privileges.
NOTE: Use the appropriate syntax for **umount** for your operating system. The above example is for Linux.
4. **Restart** the SAP ASE instances.

OR

Reboot the target/staging host

Reboot the target/staging host corresponding to the concerned dSources/VDBs. This option is probably the cleanest way of clearing the source engine's leftover dSource/VDB mounts, but should only be used if the concerned target/staging host(s) do not host any other databases that are not part of the uncontrolled failover.

Enabling Environments, dSources, and VDBs

1. Enable the source environment and then the target environment.
2. Enable the dSources and VDBs on the target engine.
3. Mounts will be moved to a new Delphix(Target Engine) and can be checked by using the following command:
mount | grep <target-engine-ip>
These mounts still use the source engine's UUID and not the target engine's UUID, but validated/manual syncs on dSources and refreshes/rewinds on VDB operations will be successful.

Related Topics

- [Uncontrolled Failover](#)

SQL Server Environment and Data Sources

Enabling Environments and dSources


1. Enable the proxy(staging) environments followed by the source environments.
2. Enable dSources on the target engine.
3. Delete the staging database stale mounts from the connector directories. An example directory: **C:\ProgramFiles\Delphix\DelphixConnector\<Old Engine ID >-staging-<X>** where **X** is an arbitrary id representing the staging database.
4. Delete the staging database from the SQL instance on the proxy host using the following convention: **<Old Engine ID>-staging-<X>** where **X** is an arbitrary id representing the staging database.

Enabling VDBs on the Standalone host

1. Delete the VDB directories stale mounts from: **C:\ProgramFiles\Delphix\DelphixConnector\<Old Engine id>-vdb-<X>** where **X** is an arbitrary id representing the virtual database.
2. From the SQL instance delete VDB databases that are in recovery pending state.
3. Enable VDBs.
New mounts will be created and a new engine ID will reflect on those mounts.


Enabling VDBs on Failover cluster as Target

1. Once the source Engine is powered down login to **Failover cluster**.
2. Go to **Windows Administrative Tools** and select **Failover cluster manager**
3. On the left pane of the **Failover cluster manager** go to the **Roles**.
As shown in the image below you will see the Failover cluster in the following state with the indicated status.




4. In order to fix the state go to the **Storage** section and select **Disks**. The **dlpx-XXXXXXXXXX** mounts attached to the Cluster Disk will be in offline states. These mounts will be equal to the total number of VDBs


provisioned on that cluster as the target.




5. Right-click on the dlpx mount, and select **Remove from SQL Server (Instance Name)**.




6. Click **Yes**.




In the **Assigned To** column, you will see **Available Storage**.




- Right-click on the **dlpx mount** and click on **Remove**.




- The **dlpx mount** will be deleted. Repeat the same process with other **dlpx mounts** assigned to **Available Storage**.
- Under **Roles**, you would see that the **SQL Server service** is now in a **Failed** state.




- In order to fix it right-click on the service and click **Stop Role**.




- Once the Stop role is executed properly, right-click on the **SQL cluster service** and click on **Start Role**.



- The SQL cluster service should now be online. Under the **Resources** tab the Storage, Server Name, Other Resources and Roles should have Online status. Ensure that the Disks and Nodes are also in an online state.



13. Make sure that both the nodes are in a healthy state.



Related Topics

- [Uncontrolled Failover](#)

Replicas and Failover

Replication recreates objects on the target system in a replica that preserves object relationships and naming on the target server without interfering with active objects on the system. Objects within a replica are read-only and disabled until a replica is failed over, at which point they can be activated. VDBs and dSources within a replica can be used as the source for provisioning new VDBs.

Replicas

A replica contains a set of replicated objects. These objects are read-only and disabled while replication is ongoing. To view replicated objects, look under **namespace** in the CLI. Or in the GUI:

1. Click **System**.
2. Select **Replication**.
3. Under **Received Replicas**, select the replica.

On this screen, you can browse the contents of replicas, as well as failover or delete individual replicas. As described in the [Replication Overview](#) topic, databases (dSources and VDBs) and environments are included within the replica.

Deleting or Failing over a Replica

Deleting or failing over a replica will sever any link with the replication source. Subsequent incremental updates will fail, requiring the source to re-establish replication. Failover should only be triggered when no further updates from the source are possible, as in a disaster scenario.

Multiple replicas can exist on the system at the same time. Active objects can exist in the system alongside replicas without interfering with replication updates. You can also use VDBs and dSources within a replica as a source when provisioning. For more information, see [Provisioning from Replicated Data Sources or VDBs](#).

Failover and Conflict Resolution

To activate the objects in a replica, you must first fail over the replica. This will sever replication and move the objects to the live system, where they can be manipulated in the same fashion as other objects on the system.

Objects that are failing over can conflict with objects in the live system. One reason for conflicts is identical names. For example, Groups will conflict if the failing over Group has the same name as a Group in the live system, as will most other objects including Environments, dSources, and VDBs. By default, active objects with conflicting names will cause an error at the time of failover. The error message will indicate which object(s) have conflicting names.

For all other object types, because the replica objects are read-only, you must rename the active objects before the failover operation can complete successfully.

Conflicting names

Given that conflicting names prevent failover from succeeding, best practices in a disaster recovery situation are to leave the target system completely passive with no active objects until the time that failover is required.

Conflicting plugins

Due to the potential for plugin version incompatibilities, if any of your replicated objects are plugin-based, we highly recommend that you leave the target system completely passive with no active objects until the time that failover is required. For more information, see the [Delphix Engine Plugin Management](#) page.

⚠ When environments are consolidated, the Delphix engine maintains the existing environment configuration and privilege elevation profile; and merges replicated objects with them. Therefore, you must check and update the configured privilege elevation profiles when Environments are consolidated.

Once a replica is failed over, the objects are active but will be automatically disabled.


Manual Conflict Resolution

To resolve conflicts in Groups, you may either rename the conflicting group on the target replication engine or if the source replication engine is still available, rename on the source replication engine and send a replication update, before failover. For all other object types, because the replica objects are read-only, you must rename the active objects on the replication source engine and send a replication update before the failover operation can complete successfully.

Automatic Conflict Resolution

Starting with 6.0.2.0, most of the object conflicts can be resolved automatically. When you select "Automatically resolve object conflicts", replica objects like Groups will be renamed whereas objects like Environments will be merged and consolidated.

As of 6.0.5.0, the Automatic Conflict Resolution option will be chosen by default.



After automatic conflict resolution and successful failover, a report is presented with a list of objects that needed conflict resolution.

 ip-10-110-198-63-5 
**Failover**

- Renamed GROUP from "Untitled" to "Untitled-ip-10-110-198-63-5"
- Consolidated UNIX_HOST "172.16.102.163" with "172.16.102.163"
- Consolidated UNIX_HOST "bbrac14" with "bbrac14"
- Consolidated UNIX_HOST "172.16.101.163" with "172.16.101.163"
- Consolidated UNIX_HOST "172.16.101.162" with "172.16.101.162"
- Consolidated UNIX_HOST "cnrac11" with "cnrac11"
- Consolidated UNIX_HOST "172.16.102.162" with "172.16.102.162"
- Consolidated WINDOWS_HOST "10.110.244.248" with "10.110.244.248"
- Consolidated WINDOWS_HOST "10.110.227.152" with "10.110.227.152"
- Consolidated WINDOWS_HOST "10.110.156.45" with "10.110.156.45"
- Consolidated WINDOWS_HOST "10.110.156.86" with "10.110.156.86"
- Consolidated WINDOWS_CLUSTER "fc8uoqar268" with "fc8uoqar268"
- Consolidated ORACLE_CLUSTER "bbrac1416" with "bbrac1416"
- Consolidated ORACLE_CLUSTER "cnrac1113" with "cnrac1113"
- Consolidated ORACLE_CLUSTER_NODE "bbrac16" with "bbrac16"
- Consolidated ORACLE_CLUSTER_NODE "bbrac14" with "bbrac14"
- Consolidated ORACLE_CLUSTER_NODE "bbrac15" with "bbrac15"
- Consolidated ORACLE_CLUSTER_NODE "cnrac13" with "cnrac13"
- Consolidated ORACLE_CLUSTER_NODE "cnrac12" with "cnrac12"
- Consolidated ORACLE_CLUSTER_NODE "cnrac11" with "cnrac11"
- Consolidated WINDOWS_CLUSTER_NODE "10-110-156-86.ad.delphix.com" with "10-110-156-86.ad.delphix.com"
- Consolidated WINDOWS_CLUSTER_NODE "10-110-156-45.ad.delphix.com" with "10-110-156-45.ad.delphix.com"

Replica Type

Replication

Replicated Objects

No replicated objects

Description**Environments**

No replicated environments

Automatic conflict resolution is not supported for objects like dSources and VDBs.

Enabling Databases and Environments

Objects may refer to states (IP addresses, mount paths, etc) that differ between the source and target system. Because of this, all databases and objects within a replica automatically start in the disabled state after a failover. This allows the administrator to alter configuration prior to enabling databases and environments, without the system inadvertently connecting to invalid systems.

After failover is complete, you must explicitly enable all dSources, VDBs, and environments. If you need to change any configuration for the target environment, you can do so prior to enabling the objects. In the event that a failing-over environment is consolidated with a live system environment, it must be refreshed before all of it's databases can be used.

Related Topics

- [Replication Overview](#)
- [Provisioning from Replicated Data Sources or VDBs](#)

Provisioning from Replicated Data Sources or VDBs

This topic describes how to provision from a replicated dSource or VDB. The process for provisioning from replicated objects is the same as the typical VDB provisioning process except for the need to first select the namespace containing the replicated object.

Prerequisites

You must have done the following:

- replicated a dSource or a VDB to the target host, as described in [Replication Overview](#).
- added a compatible target environment on the target host as described in [Provisioning Oracle VDBs: An Overview](#).
- installed on the target host any App Data plugin which the replicated objects depend on

Procedure

1. Login to the **Delphix Management** application for the target host.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **replica** that contains the dSource or VDB to be provisioned.
5. The provisioning process is now identical to the process for provisioning standard objects.

 If the dSource or VDB belongs to a plugin, the plugin must exist on the target engine to be successfully provisioned. The target's plugin's version must also be equal to or higher than the source's plugin's version. For more information see the [Delphix Engine Plugin Management](#) topic.

Post-Requisites

Once the provisioning job has started, the user interface will automatically display the new VDB in the live system.

Related Topics

- [Replication Overview](#)
- [Provisioning Oracle VDBs: An Overview](#)

Replication User Interface

The replication user interface helps users manage replication on both the source and the target. Replication consists of a profile-replica pair. You can view and edit the replication profile on the source engine, and view the replica on the target engine.

Sources for Replication

The **Replication Profiles** section provides you with several configuration options for your replication profiles. This makes it possible to replicate objects from a single source to multiple targets. Each profile defines the set of data objects and the associated configuration between a single source and target.

Targets for Replication

The **Received Replicas** section shows the set of all objects in the replica and allows you to initiate failover.

Canceling a Replication Job


You must cancel the replication job at the source engine. Canceling the replication job at the target engine results in the re-starting of the receiving job at the target engine. If the source engine is down, then you can cancel the replication job at the target engine. If you re-start the source engine where the replication job was previously running, the source engine will attempt to restart the replication send job.

Replication UI

The Replication UI consists of the following sections.

- **Replication Profiles**

The screenshot and descriptions below illustrate the capabilities in the **Replication Profiles** section.



1	<p>Create Profile Button</p> <p>It allows you to configure a Replication Profile.</p>								
2	<p>Replication Profiles List</p> <p>Provides a list of existing replication profiles. Click a profile in this list to view its details.</p>								
3	<p>Received Replicas List</p> <p>Provides a list of all existing replicas on this Delphix Engine. Click a replica in this list to view its details.</p>								
4	<p>Status Box</p> <p>Shows the replication status of the selected profile, including:</p> <ul style="list-style-type: none"> • The result of the most recent or current replication event • Statistics for the replication run such as data transferred, duration, and average throughput <p>In the upper left-hand corner, an icon summarizes the replication status. There are four possible status icons:</p> <table border="1" data-bbox="531 1005 1423 1400"> <tr> <td data-bbox="531 1005 778 1122"></td> <td data-bbox="778 1005 1423 1122">This icon appears while a replication job is in progress.</td> </tr> <tr> <td data-bbox="531 1122 778 1205"></td> <td data-bbox="778 1122 1423 1205">This icon appears after a successful replication job.</td> </tr> <tr> <td data-bbox="531 1205 778 1288"></td> <td data-bbox="778 1205 1423 1288">This icon appears when a replication job has failed.</td> </tr> <tr> <td data-bbox="531 1288 778 1400"></td> <td data-bbox="778 1288 1423 1400">This icon appears when a replication job was canceled.</td> </tr> </table>		This icon appears while a replication job is in progress.		This icon appears after a successful replication job.		This icon appears when a replication job has failed.		This icon appears when a replication job was canceled.
	This icon appears while a replication job is in progress.								
	This icon appears after a successful replication job.								
	This icon appears when a replication job has failed.								
	This icon appears when a replication job was canceled.								
5	<p>Configuration Options</p> <p>Configuration Options for the selected replication profile. These include:</p> <ul style="list-style-type: none"> • Description – Free text description of the profile • Target Engine – The Delphix Engine on the receiving end of this replication pair • Automatic Replication – If enabled, shows the frequency and time that regular replication will be run • Traffic options – Summarizes the traffic options with which this profile has been configured 								

<p>6</p>	<p>Object Selection Tree</p> <p>Shows all of the objects, such as groups, dSources, VDBs, and Self-service (JetStream) data layouts, that you have selected for replication in this replication profile. If you select Entire Delphix Engine, all objects on the engine will be replicated, and thus the tree is collapsed.</p>
<p>7</p>	<p>Replicate Now Button</p> <p>Begins the replication process</p>
<p>8</p>	<p>Delete Button</p> <p>Allows you to delete the current profile</p>

• Received Replicas

The screenshot and descriptions below provide more details of the functionality of this section.

The screenshot displays the 'Received Replicas' section of the Delphix Engine interface. On the left, a sidebar contains 'Replication Profiles' and 'Received Replicas'. The main area shows a replica for 'ip-10-110-196-234-1'. A status box (1) indicates 'Last Replication Successful' with a timestamp and duration. Below this, the 'Replica Type' (2) is set to 'Replication'. The 'Environments' section lists 'target' and 'source'. The 'Users' section lists 'admin'. The 'Replicated Objects' tree (3) shows a hierarchy starting with 'dSources and VDBs' and an 'Untitled' folder containing several objects. A 'Failover Now' button (4) is visible in the top right, and a trash icon (5) is next to it.


<p>1</p>	<p>Status Box</p> <p>Similar to the Replication Profile status box, this shows the most up-to-date status information for the replica on the target</p>
<p>2</p>	<p>Replicated Environments</p> <p>Replicating a dSource or VDB will automatically replicate any environments associated with those objects. For more information, see Replication Overview.</p>
<p>3</p>	<p>Replicated Objects Tree</p> <p>A read-only view of the objects in this replica</p>

	<p>Failover Button</p> <p>Initiates a failover for this replica</p>
	<p>Delete Button</p> <p>Deletes this Replica on the target. This does not have an effect on the corresponding profile on the source engine.</p>

Create New Replication Profile Wizard

Perform the following steps to create a new replication profile.

1. Login to the **Delphix Management** application.
2. Navigate to **System > Replication**.



3. Next to **Replication Profiles**, Click the **plus** icon. The **Add Replication Profile** wizard screen appears.
4. Enter the **Name** of the replication profile and an optional **Description**, and click **Next**.
5. To specify the **Object**, select the **Profile Type** as **Replication Profile** and select from the **Objects Being Replicated** to a target engine. Click **Next**.

(i) Selected Objects

- Some selected objects may have dependencies – other objects that will be pulled into replication because they share data. Objects that will be replicated are confirmed with a blue chain link icon.
 - **Note that this is not guaranteed to be the full set of dependent objects, but rather is the best guess. The full set of objects and their dependents will be calculated at the time of replication.**
- When selecting objects, you can select the entire server (**Entire Delphix Engine**) or a set of groups, dSources, VDBs, and Jet Stream data layouts.
- When replicating a group, all dSources and VDBs currently in the group, or added to the group at a later time, will be included.
- If you select a Delphix Self-Service data template, all data containers created from that template will be included. Likewise, if you select a data container, its parent data template will be included.

- If you select the entire server, all groups and Delphix Self-Service objects will be included.
- Regardless of whether you select a VDB individually or as part of a group, the parent dSource or VDB (and any parents in its lineage) are automatically included. This is required because VDBs share data with their parent object. In addition, any environments containing database instances used as part of a replicated dSource or VDB are included as well.
- When replicating individual VDBs, only those database instances and repositories required to represent the replicated VDBs are included. Other database instances that may be part of the environment, such as those for other VDBs, are not included.
- Non-data objects (Delphix users, roles, permissions, authorizations, policies, database configuration templates) that are associated with selected objects will be automatically included during replication. They will not be shown on the selection interface.

6. For **Target Engine**, enter the **hostname** or **IP address** for the target Delphix Engine. Enter the **username** and **password** of a user who has Delphix Admin-level credentials on the target Delphix Engine. If the username or password changes on the target Delphix Engine, you must update these settings on the source Delphix Engine. Click **Next**.
7. By default, automatic replication is disabled and you must trigger replication updates manually. To **Schedule** automatic replication, click the **Enabled** checkbox. In the **Automatic Replication** field, enter the frequency and starting time for replication updates to the target Delphix Engine. Once you have entered and saved your replication settings, you will also see an option to trigger replication immediately with the **Replicate Now** button. Under **Traffic Options**, select whether you want to **Encrypt** traffic or **Limit bandwidth** during replication updates. Click **Next**.
8. The final summary tab will enable you to review your configurations. Click the **Back** button to go back and to change any of the configurations or click **Submit** to complete the wizard.

Viewing and Editing an Existing Replication Profile

Perform the following steps to view and edit existing replicas.

1. Login to the **Delphix Management** application.
2. Navigate to **System > Replication**.
3. Under the **Received Replicas** section, select a replica.
4. You can edit the **Name** and **Description** fields. All other fields are view-only.

Perform the following steps to edit the fields.

1. Click the **pencil** icon next to the corresponding field or group of fields to edit the fields.
2. To save the edits and/or selections, click the **Checkmark** icon.
3. To cancel the edits and/or selections, click the **cross** icon.

Related Topics

- [CLI Cookbook: Replication](#)
- [Configuring Replication](#)
- [Replicas and Failover](#)

Selective Data Distribution

These topics describe concepts and procedures for using Selective Data Distribution (SDD) to replicate data from one Delphix Engine to another.

- [Support Matrix for Selective Data Distribution](#)
- [Selective Data Distribution Overview](#)
- [Selective Data Distribution Use Cases](#)
- [Selective Data Distribution User Interface](#)
- [Configuring Selective Data Distribution](#)
- [Selective Data Distribution and Failover](#)

Support Matrix for Selective Data Distribution

SDD is supported across all environments supported by each data source. Refer to the support matrix for each data source below.

- [Oracle Matrix](#)
- [SQL Server Matrix](#)
- [SAP ASE Matrix](#)
- [Db2 Matrix](#)
- [EBS Matrix](#)
- [HANA Matrix](#)
- [PostgreSQL Matrix](#)

Selective Data Distribution Overview

The selective data distribution technology permits the distribution of masked data between Delphix Engines without bringing over the unmasked parent source. These engines must be the identical versions of the Delphix Data as a Service Engine with the Delphix Masking Engine. Otherwise, they can be asymmetric in terms of engine configuration. You can provision VDBs from distributed masked objects, allowing for the geographical distribution of data and remote provisioning.

You can run selective data distribution ad hoc, but it is typically run according to a predefined schedule. After the initial update, each subsequent update sends only the changes incurred since the previous update. Selective data distribution does not provide synchronous semantics, meaning that the data distributed to the target is only as current as of the most recent update.

Selective Data Distribution Features

As virtual appliances, you can backup, restore, replicate, and migrate data objects between Delphix Engines using features of VMWare and the underlying storage infrastructure. In addition to the replication capabilities provided by this infrastructure, selective data distribution permits the distribution of masked data between Delphix Engines. The sources received on a target Delphix Engine do not include the original parent source, thereby making the original source inaccessible from the target.

Selective data distribution is configured on the source Delphix Engine. It first copies a subset of masked VDBs to a target Delphix Engine, then sends incremental updates either manually or according to a schedule. As illustrated below, sensitive data from the Production Data Center is brought into Delphix as a dSource. The Masking Engine then masks the dSource data as a VDB. Synchronously, DxFS redacts the sensitive data within that dSource before sending it across into the Non-Production datacenter using Delphix Replication. Using SDD capability, sensitive data is never exposed because it is protected both at the dSource and VDB layers.



You can use replicated masked VDBs to provision new VDBs on the target Delphix Engine. The provisioned VDBs contain the data in their masked parent and are therefore also considered masked. You can refresh these VDBs to snapshots sent as part of an incremental replication update, as long as you do not destroy the parent object on the replication source. For more information, see [Provisioning from a Replicated Data Sources or VDBs](#).

During replication, replicated masked VDBs are maintained in an alternate replica and are not active on the target side. The failover of a selective data distribution replica is not supported.

Selective Data Distribution Details

When you select masked objects for selective data distribution, the engine will automatically include any dependencies, such as environments, associated with the VDB. The parent dSource and any parent VDBs are not included automatically. The data associated with parent objects are selectively included for disk space efficiency, but data in the parent dSource and VDBs that the masked VDB does not need are excluded.

During replication, the Delphix Engine will negotiate an SSL connection with its server peer to use SSL_RSA_WITH_RC4_128_MD5 as the cipher suite, and TLSv1 as the protocol.

Only database objects and their dependencies are copied as part of a selective data distribution operation, including:

- Masked VDBs
- Environments
- Environment configuration (users, database instances, and installations)

The following objects are not copied as part of a selective data distribution operation:

- Parent dSources of masked VDBs - The storage blocks for non-sensitive dSource data are sent from the source to the target replication host. This storage is displayed under Held Space, for more information see [An Overview of Held Space](#).
- Groups of the parent dSources
- Users and roles
- Policies
- VDB (init.ora) configuration templates
- Events and faults
- Job history
- System services settings, such as SMTP

Resumable Selective Data Distribution

A single selective data distribution instance can fail for a number of environmental and internal reasons. However, using the Resume feature, you can restart selective data distribution from an intermediate point; no data is retransmitted. Selective data distribution is resumable across machine reboot, stack restart, and network partitions. The resumable replication feature is fully automated and does not require or allow any user intervention.

For example, you can resume a large, time-consuming initial distribution or incremental update after it is interrupted. Suppose a selective data distribution profile has already been configured from a source to a target. A large, full send from the source begins that is expected to take weeks to complete. Halfway through, a power outage at the data center that houses the source causes the source machine to go down and only come back up after a few hours. On startup, the source will detect that a selective data distribution was ongoing, automatically re-contact the target, and resume the distribution where it left off. In the user interface (UI) on the source, the same selective data distribution send job will appear as active and continue to update its progress. However, in the UI of the target, a new distribution receives job will appear, although it will track its progress as a percentage of the entire replication.

Selective data distribution will not resume after failures that leave the source and target connected. For example, if a storage failure on the target, such as an out-of-space error, causes a distribution to fail, then the source and target remain connected. As a result, the Engine will discard state data associated with the failed Replication operation.

Selective Data Distribution Restrictions

- Only masked VDBs can be added to a selective data distribution spec. You cannot add dSources, groups, or the entire domain.
- Only masked VDBs with a Snapshot Policy of None should be added to a selective data distribution spec.
- Unmasked VDBs cannot be added to a selective data distribution spec.
- VDBs that undergo selective data distribution and their children cannot be selectively redistributed to another target.
- You cannot go to the target engine and create a selective data distribution spec that includes VDBs that are present because of selective data distribution. However, you can replicate this data using a traditional replication spec.

Best practices for using Selective Data Distribution are described in the [Selective Data Distribution Best Practices](#) knowledge base article.

Selective Data Distribution Supported Platforms

Selective Data Distribution supports the following platforms:

- Oracle
- Microsoft SQL Server
- SAP ASE (Sybase)

- [Db2](#)

Related Topics

- [Selective Data Distribution](#)


Selective Data Distribution Use Cases

Selective data distribution is a flexible tool that allows you to move dSources and virtual databases (VDBs) between Delphix Engines. Selective data distribution relies on the replication infrastructure to transmit masked data, as such selectively distributed data has much of the same capabilities as replicated data. These topics describe several different use cases for selective data distribution.

Geographically Distributed Development

The Delphix Engine allows you to provision VDBs from distributed dSources and VDBs, as described in [Provisioning from Replicated Data Sources or VDBs](#). This allows dSources that are linked in a single central location to be geographically distributed so that developers can provision VDBs remotely without having to sync from the source database in multiple locations. Selective distribution ensures that masked VDBs are sent to the target without transmitting the original unmasked source. As such, the original unmasked source will not be accessible on the target.

Selective data distribution does not support failover. In this environment, the link between the source and the target is never broken, except when remote VDBs need to be preserved. You can refresh remote VDBs as long as the parent objects continue to exist on the source. If they are deleted, then remote VDBs will continue to function but cannot be refreshed.



Because there is no failover, this technology can support more complex topologies such as 1-to-many and many-to-1. Chained distribution (replicating from Site A -> Site B -> Site C) is not supported.

Best Practices

For geographical distribution, follow these best practices:

- Because each replication stream induces load on the source system:
 - Minimize the number of simultaneous replication updates
 - If possible, avoid heavy VDB workloads on the source
- On the target, the provision only from sources that are effectively permanent. If a source is deleted remote VDBs can no longer be refreshed.
- Provision additional storage capacity on the target
 - Remotely provisioned VDBs can consume shared storage on the target even when the parent is deleted on the source

Migration



You cannot use selective data distribution for data migration. Full replication is needed for data migration.

Related Topics

- [Selective Data Distribution](#)


Selective Data Distribution User Interface

Sources for Selective Data Distribution

The **Replication Profile** section continues to handle Replication functions, but also allows for the configuration of **Selective Data Distribution (SDD)** profiles. This makes it possible to distribute objects from a single source to multiple targets. Each profile defines the set of data objects and the associated configuration between a single source and target.

Replication/Selective Data Distribution Section

This section is used to illustrate the capabilities in the **SDD Profiles** section. From the Delphix Management portal, navigate to **System > Replication**.



Select the plus icon to add a profile, this will open a new window for configurations. Below the **Replication Profiles** header and the plus icon is a list of existing Replication Profiles and SDD profiles. Select a profile in this list to view its details.

The **Received Replicas List** shows all replicas received to this Delphix Engine, including normal Replication and SDD replicas. Click a **replica** in this list to view its details.

A **Status Box** appears when a profile is selected from the list that shows the distribution status of the selected profile. This includes the result of the most recent or current distribution event and statistics for the distribution run, such as data transferred, duration, and average throughput. The **Profile Type** shows the type of the selected profile or replica.

There are several configuration options available for the selected profile or replica listed below the Profile Type. These sections, in addition to the profile name itself, can be modified by selecting the pencil icon.

- Description – Free text description of the profile
- Target Engine – The Delphix Engine on the receiving end of this data distribution pair
- Automatic Replication – If enabled, shows the frequency and time that regular distribution will be run
- Traffic Options – Summarizes the traffic options with which this profile has been configured

The **Objects Being Replicated** shows all of the masked objects selected for distribution in this SDD profile. The **Replicate Now** button initiates the distribution process and the **bin icon** to the left allows for current profile deletion.

The screenshot displays the 'Replication Profiles' section in the Delphix Management console. The profile 'sdd' is selected, and its configuration is shown on the right. A 'Replicate Now' button is visible in the top right corner. The configuration includes:

- Last Replication Successful:** Last Successful Replication Jul 9, 2020 4:27:18 PM (2 minutes ago). Duration: 0:01:04 | 1.71GB | 46.99MB/s.
- Profile Type:** Selective Data Distribution.
- Description:** (Editable)
- Target Engine:** 6021.dlpxdc.co (Editable). Note: = will be replicated.
- Automatic Replication:** Disabled (Editable).
- Traffic Options:** None (Editable).
- Objects Being Replicated:** dSources and VDBs > Untitled > VPDB_8NQ (Source: PDB1. Masked parent: VPDB_8NQ).

Below the configuration, there is a 'Received Replicas' section with a sub-section for 'ip-10-110-227-8-1'.

Create New Profile

1. In the left-hand navigation section, click the **+** icon.
2. Enter the **name** of the SDD profile and an optional **description**.
3. For **Type**, select **Selective Data Distribution**.
4. For **Target Engine**, enter the **hostname** or **IP address** for the target Delphix Engine.
5. Enter the **username** and **password** of a user who has Delphix Admin-level credentials on the target Delphix Engine. If the username and password change on the target Delphix Engine, these settings must be updated on the source Delphix Engine.
6. Automatic replication is disabled by default, which means the distribution updates must be triggered manually via the **Replicate Now** button. To enable automatic distribution, click the **Enabled** checkbox.
7. In the **Automatic Replication** field (if enabled), enter the **Frequency** and **Starting Time** for distribution updates to the target Delphix Engine.

! Automatic replication uses Quartz for scheduling. Starting with Delphix version 4.2, the Quartz-formatted string is editable via the **Advanced** option. An example **Cron String** would be seen as: `0 0 * * ?`

8. Under **Traffic Options**, select to **Encrypt** traffic or **Limit bandwidth** during distribution updates.

9. In the right-hand column under **Objects Being Replicated**, select the **boxes** next to the objects to distribute. Only masked VDBs can be used and multiple masked VDBs cannot be added to the same profile.
10. Click **Create Profile** to submit the new profile. This saves the SDD profile details. Leaving the **Create** page prior to submitting the profile will result in the draft being discarded.

Related Topics

- [Selective Data Distribution](#)

Configuring Selective Data Distribution

This topic describes how to configure Selective Data Distribution between Delphix Engines. Selective data distribution relies on the replication infrastructure to transmit masked data, as such configuring selective data distribution is similar to configuring replication.

Prerequisites

- The replication target can be on the same or newer version than the replication source. Starting 6.0.10.0, replication operation can not go beyond engine versions released more than 12 months apart. For more information, see [Replication Overview](#).
- The target Delphix Engine must be reachable from the source engine.
- The target Delphix Engine must have sufficient free storage to receive the replicated data.
- The user must have administrative privileges on the source and the target engines.

Configuring the Network

Replication and selective data distribution operate using a private network protocol between two Delphix Engines. Apart from standard network considerations for performance, no additional configuration is required for replication. Replication and selective data distribution can run over dedicated networks by configuring routing to direct traffic destined for the target IP address over a specific interface. The selective data distribution process can recover from transient network outages, but extended outages may cause the process to start from the previous update.

The selective data distribution network protocol uses TCP port 8415. If there is a firewall between the source and target that is blocking this port, then there are two possible solutions:

- Enable port 8415 on the firewall in order to allow connections to this port from the source to the target.

or

- Selective data distribution can connect through a SOCKS proxy if one exists. Configure the SOCKS proxy address and port by connecting to the command-line interface (CLI) as a system administrator and navigating to "service proxy" to update the socks configuration. Example:

```
dlpx-engine> service proxy
dlpx-engine service proxy> update
dlpx-engine service proxy update *> set socks.enabled=true
dlpx-engine service proxy update *> set socks.host=10.2.3.4
dlpx-engine service proxy update *> set socks.username=someuser
dlpx-engine service proxy update *> set socks.password=somepassword
dlpx-engine service proxy update *> commit
dlpx-engine service proxy> get
  type: ProxyService
  https:
    type: ProxyConfiguration
    enabled: false
    host: (unset)
    password: (unset)
    port: 8080
    username: (unset)
  socks:
```

```

type: ProxyConfiguration
enabled: true
host: 10.2.3.4
password: *****
port: 1080
username: someuser

```

Note that SOCKS port 1080 is used by default, but you can override it.

Configuring the Source Delphix Engine

1. On the source Delphix Engine, click **System**.
2. Select **Replication**.
3. In the left-hand navigation section, click **Create Profile**.
4. Enter the **name** of the replication profile and an optional **description**.
5. Select type **Selective Data Distribution**.
6. For **Target Engine**, enter the **hostname** or **IP address** for the target Delphix Engine.
7. Enter the **username** and **password** of a user who has Delphix Admin-level credentials on the target Delphix Engine. If the username and password change on the target Delphix Engine, you must update these settings on the source Delphix Engine.
8. By default, automatic replication is disabled, meaning that you must trigger replication updates manually. To enable automatic replication, click the **Enabled** checkbox.
9. In the **Automatic Replication** field, enter the **Frequency** and **Starting Time** for replication updates to the target Delphix Engine. Once you have entered and saved your replication settings, you will also see an option to trigger replication immediately with the **Replicate Now** button.

Quartz scheduling

Automatic replication uses Quartz for scheduling. Starting with Delphix version 4.2, the Quartz-formatted string is editable via the **Advanced** option.

10. Under **Traffic Options**, select whether you want to **Encrypt** traffic or **Limit bandwidth** during replication updates.

Encrypting Traffic

By default, replication streams are sent unencrypted. This provides maximum performance on a secure network. If the network is insecure, encryption can be enabled. Note that encrypting the replication stream will consume additional CPU resources and may limit the maximum bandwidth that can be achieved.

Limiting Bandwidth

By default, replication will run at the maximum speed permitted by the underlying infrastructure. In some cases, particularly when a shared network is being used, replication can increase resource contention and may impact the performance of other operations. This option allows you to specify the maximum bandwidth that replication can consume.

11. In the right-hand column, under **Objects Being Replicated**, click the **checkboxes** next to the objects you want to replicate.

Selected Objects

- You can only select masked VDBs for selective data distribution.

- The parent dSource or VDB (and any parents in its lineage) are NOT automatically included. Some of the data from the parent may be included for disk space optimization. In addition, any environments containing database instances used as part of a replicated VDB are included as well.
- When replicating individual VDBs, only those database instances and repositories required to represent the replicated VDBs are included. Other database instances that may be part of the environment, such as those for other VDBs, are not included.

12. Click **Create Profile** to submit the new profile. This saves the replication profile details. If you leave the **Create** page prior to submitting the profile, the draft replication profile will be discarded.

Configuring Replication and Multiple Target Engines through the CLI

You can also configure replication on the Source Delphix Engine by using the replication spec in the command line interface. For more information, see the topics under [CLI Cookbook: Replication](#).

Enabling Configuration of Multiple Replication Profiles

Learn how to configure and use functionality for multiple replication profiles on the source using the replication profiles in the [Replication User Interface](#).

Configuring the Target Delphix Engine

No additional configuration on the target is needed. Selectively distributed objects will appear in an alternate replica that mirrors the original object layout. To view these replicas:

1. Click **System**.
2. Select **Replication**.
3. Review items listed under **Received Replicas**.

Alternatively, you can view replicas under **namespace** in the command-line interface (CLI). All replicated objects are read-only. For more information about managing replicas, see [Selective Data Distribution and Failover](#).

Multiple sources can replicate to the same target, allowing for the flexible geographical distribution of data. You can create and manage objects on the target server without affecting subsequent updates. However, you cannot use selective data distribution for disaster recovery.

Related Topics

- [Selective Data Distribution](#)

Selective Data Distribution and Failover

Selective data distribution recreates objects on the target system in a replica that preserves the object relationships and naming on the target server without interfering with active objects on the system. Objects within a replica are read-only and disabled. They cannot be failed over. However, you can use virtual databases (VDBs) within a replica as the source for provisioning new VDBs.

Replicas

A replica contains a set of replicated objects. These objects are read-only and disabled while replication is ongoing. To view replicated objects:

1. Click **System**.
2. Select **Replication**.
3. Under **Received Replicas**, select the replica. On this screen, you can browse the contents of replicas or delete individual replicas. As described in the [Selective Data Distribution Overview](#), VDBs and environments are included within the replica.

You can also view replicated objects under **namespace** in the command-line interface (CLI).

 Deleting a replica will sever the link with the replication source. Subsequent incremental updates will fail, requiring the source to re-establish replication.

Multiple replicas can exist on the system at the same time. Active objects can exist in the system alongside replicas without interfering with replication updates. VDBs within a replica can also be used as a source when provisioning. For more information, see [Provisioning from Replicated Data Sources or VDBs](#).

Failover

Selective data distribution does not support the failover of replicas.

Related Topics

- [Selective Data Distribution](#)
- [Selective Data Distribution Overview](#)
- [Provisioning from Replicated Data Sources or VDBs](#)

Delphix Self-Service

 Unknown macro: 'navbox'

Delphix Self-Service Admin Guide

This section covers the following topics:

- [Getting Started with Delphix Self-Service](#)
- [Delphix Self-Service Concepts](#)
- [Navigating the Delphix Self-Service Admin Interface](#)
- [Understanding Data Templates](#)
- [Understanding How to Manage Data Template Details](#)
- [Understanding Data Containers](#)
- [Delphix Self-Service Data Container Activities](#)
- [Using Masked Data Sources with Delphix Self-Service](#)
- [Understanding Delphix Self-Service User Management](#)
- [Working with Multiple Container Owners](#)
- [Understanding Bookmarks](#)
- [Understanding Delphix Self-Service Usage Management](#)

Getting Started with Delphix Self-Service

Welcome to Delphix Self-Service

The Delphix Engine has greatly improved the speed at which end users can get the data that they need. While end-users reap the benefits, they do not typically interact with the Delphix Engine directly, nor are they necessarily even aware that they are using it. End users most commonly file tickets for data management operations and wait for the tickets to be serviced by their IT organization. Delphix data management workflows allow database administrators (DBAs) to respond to these tickets much more quickly and reliably, but DBAs are often overloaded, and resolving high-priority issues takes precedence over requests from users. Requiring interactions between users and IT for every data operation is inefficient and can lead to unwanted delays.

The goal of Delphix Self-Service is to create a clear separation of IT infrastructure and data management. As with the current Delphix Engine, IT administrators and DBAs continue to control decisions about how resources such as virtual databases (VDBs) and vFiles are allocated. However, with Delphix Self-Service, administrators can also assign these resources directly to a user. A user has the ability to control what data their container presents, even though the details of the physical resources are hidden from them. This separation of roles empowers Delphix Self-Service users to get the data they need, when they need it while providing administrators with the controls to ensure resources are accounted for appropriately.

User Roles and Permissions

Delphix Self-Service has two types of users:

Admin User

Admin users have full access to all report data and can configure Delphix Self-Service, additionally, they can:


- use the Delphix Engine to add/delete users
- change tunable settings
- add/delete tags
- create and assign data templates and containers

Data User

Data users have access to production data provided in a data container. The data container provides these users with a playground in which to work with data using the **Self-Service Toolbar**.

Login

1. Access Delphix Self-Service by opening a web browser and using the **IP address** or **DNS qualified hostname**.
2. Login with the **Delphix Admin User ID** and **Password** provided for you.
3. Upon successful login, you should be able to click on your username at the top right corner and select a Self-service view from the dropdown menu. This will open the Delphix Self-Service interface.



Related Topics

- [Delphix Self-Service Admin Guide](#)

Delphix Self-Service Concepts

Data Sources

A data source can be a database, an application, or a set of unstructured files. Delphix administrators configure the Delphix Engine to link to data sources, which pulls in the data of these sources. The Delphix Engine will periodically pull in new changes to the data, based on a specific policy. This, in turn, begins building a custom timeline for each data source. Additionally, the Delphix Engine can rapidly provision new data sources that are space-efficient copies, allowing users to work in parallel without impacting each other.

Data Templates

Data templates are the backbone of data containers. They are created by the Delphix administrator and consist of the data sources that users need in order to manage their data playground and their testing and/or development environments. Data templates serve as the parent for a set of data containers that the administrator assigns to users. Additionally, data templates enforce the boundaries for how data is shared. Data can only be shared directly with other users whose containers were created from the same parent data template.

Data Containers

A data container allows data users to access and manage their data in powerful ways. Their data can consist of application binaries, supporting information, and even the entire database(s) that underlie it.

A data container allows users to:


- Undo any changes to their application data in seconds or minutes
- Have immediate access to any version of their data over the course of their project
- Share their data with other people on their team, without needing to relinquish control of their own container
- Refresh their data from production data without waiting for an overworked DBA

A data container consists of one or more data sources, such as databases, application binaries, or other application data. The user controls the data made available by these data sources. Just like data sources in a template, changes that the user makes will be tracked, providing the user with their own data history.

The **Data Container Interface** lets users view the details and status of their data container and its associated data sources, as well as manipulating which data is in those sources. The **Data Container Interface** includes a section called the **Data Container Report Panel**, which displays details about each source, including the connection information needed to access it – for example, the java database connectivity (JDBC) string for a database. This connection of information is persistent and stable for the life of the data container, regardless of what data the resources are hosting.

Delphix Self-Service Data Flow

The Delphix Self-Service data flow diagram below demonstrates how a Delphix Self-Service data user accesses data sources. Data sources are connected to a Delphix Engine, which is controlled by the Delphix administrator. The Delphix administrator will connect all data sources that developers and quality assurance (QA) teams need to a data template. This data template acts as a parent source to create the data containers that the administrator will assign to data users. Data sources flow from the Delphix Engine into a data template and downstream into a data container, where a data user or users will use the data sources to complete tasks. The data container acts as a self-contained testing environment and playground for the data user. Additionally, data users are able to set, bookmark, and share data points in their container with other data users of other data containers, as long as all the data containers were created from the same parent data template.



Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Getting Started with Delphix Self-Service](#)

Navigating the Delphix Self-Service Admin Interface

The following screenshots provide a roadmap for how to navigate the primary screens and places a user will go within the Delphix Self-Service application. The application includes screens such as the **Management Overview**, **Data Platform Management**, **Users and Permissions**, and **Data User Management**.



Overview Screen

The **Overview** screen is the homepage for Delphix Self-Service. To reach this page from the Management application as an admin user, select **Self-Service** in the user login drop-down.



Data Management Overview Page

On this page, you can:



You can delete one or more templates at one time by selecting the checkboxes and clicking the Delete button.



Replicated Templates

You cannot delete replicated templates or templates that have containers.



From this column, you can see which replica the template belongs to if it is a replicated template. You can sort by this column and see all templates from the same replica.



Replica Names

This column lists replica names, not replication sources. You can view a list of replicas in the Management application at **System > Replication > Received Replicas**.

3

The number of containers associated with the template.

4

Shows the last time the template was updated.

5

Create a new data template.

Users Screen

By selecting the Users Button you can view users, whom you can then assign to data containers that you create from existing data templates.

Name	Containers	Role	Email Address
admin	2	Engine Admin	webtester1@smtptest.com
batgirl	1	Self-Service User	batgirl@mycompany.com
tom	1	Self-Service User	tom@abc.com

1

List of all configured users. You can click on a user to see their details, such as what containers they are assigned to.

2

The number of containers assigned to the user.

3

The Role column indicates whether a user is a **Jet Stream Only** user or a **Delphix Admin**.

4


Lists the associated email address for the user.

5

Select this link to add a new user or manage users.

Data Template Management

The **Data Template Management** page contains a view panel of 6 tiles on the left-hand side of the screen. Each tile reports on a variety of useful information, such as user activity, data sources, data capacity, specific details about data containers, and data templates. They also help you navigate to areas where you can complete specific tasks, such as creating a new template or container, working with data timeflows, assigning users to containers, and bookmarking important points in time.



User Roles in Admin App

To add or manage user roles and permissions, select the **Add or manage users** link from the **Users Overview** screen.

Overview

DATA TEMPLATES USERS [Add or manage users](#)

Name	Containers	Role	Email Address
admin	2	Engine Admin	webtester1@smtptest.com
batgirl	1	Self-Service User	batgirl@mycompany.com
tom	1	Self-Service User	tom@abc.com

The **Users** screen in the Management application will open. From this screen, you can add a new user, edit a user, suspend a user, or delete a user.

Users

Y + / > < ©

Username ^	Email	Type	First Name	Last Name	Work Phone N...	Mobile Phone ...	Home Phone N...
admin	webtester1@smtptes...	Admin					
batgirl	batgirl@mycompany...	Self-Service Only					
tom	tom@abc.com	Self-Service Only					

Selecting the **+** icon opens the **Add User** window. Use this window to add a user and select the user role and permissions.

Add User

Data Operations Interface for Delphix Administrators

The **Overview** page is the only interface to which Delphix Self-Service data users have access and with which they interact. The user interface is the environment in which a data user works with data in an assigned data container, using data sources from a data template.

As shown below, the **Data Operations** screen contains a list of available templates. Select a template to view details.

Template	Container	Locked By	Last Updated
Template 1	Container 1		Tue, Nov 6, 2018 10:04:42 AM
Template 2	Container 2		Tue, Nov 6, 2018 10:07:24 AM
Template 2	Container 3		Tue, Nov 6, 2018 10:09:36 AM


The following screen allows users to complete tasks using self-service operations.



For more details about how to use this interface, refer to the [Delphix Self-Service Data User Guide](#).

Administrator Settings

You can open the **Administrator Settings** dialog from the user drop-down menu. Here you can edit various settings that apply globally to all Delphix Self-Service users.



Default Bookmark Expiration

You can set a value that controls the default expiration time, in days, for Bookmarks. This setting only applies to new bookmarks that are created through the Delphix Self-Service application, not the CLI or API. Note that this only controls the default selection; users can still disable expiration or pick a different date for a specific bookmark if they wish. This setting is disabled by default.

Automatic Retry for Data Operations

To make operations on data containers more robust, Delphix Self-Service supports automatically retrying failed sources during data operations. You can specify a maximum number of retry attempts so that if an operation fails on any individual data source within a data container, it will be automatically retried until it succeeds, or until the retry limit is reached.

Automatic retry applies to any Delphix Self-Service operation that changes the data in the data container, such as Refresh, Restore, Reset, or Create Branch. This setting can be especially useful in scenarios where there are a large number of sources in a data container, and some sources fail to update the first time. If the reason for the failures was intermittent, an automatic retry may allow the sources that failed to succeed, and the operation can still complete successfully. The default number of retry attempts is **1**.

To change the number of automatic retries:

1. Click the user icon in the upper right-hand corner of the screen.
2. From the drop-down menu, select **Settings**.
3. In the field **Automatically retry data sources...**, enter a new maximum; alternatively, use the arrows to increase or decrease the number.
4. Click **Save**.

Related Topics

- [Delphix Self-Service Admin Guide](#)

Understanding Data Templates

Data Templates: An Overview

A data template represents a collection of data sources that you can provision to a user. A data source can be a dSource, VDB, or vFiles. These sources can be used in multiple data templates. Once you have created a data template, the set of data sources associated with it is fixed; you cannot add data sources to an existing template, nor can you remove data sources from it. In addition to data sources, you can define the set of metadata that is relevant for a given template – for example, notes, descriptions, names for sources that are relevant to an end-user, and other configuration details. Once you have created the template, it provides a stencil for provisioning data containers. This, in turn, enables users to have self-service access to a space-efficient copy of the data sources defined in the data template.


Data Template Activities

Data templates are managed by a Delphix admin. The admin can provision data containers from the data template and assign a data container to an end-user. The admin can also create bookmarks on the data template timeline in order to mark meaningful points in time.

When creating a data template, it is important to consider the set of users who will own data containers provisioned from it. In Delphix Self-Service, templates effectively define the boundaries of the data that users can share directly with each other. Only owners of data containers created from the same data template are able to share data using bookmarks.

Creating a Data Template and Adding Data Sources

A data template consists of an arbitrary set of dSources, virtual databases (VDBs), and vFiles. These are created and managed in the **Delphix Management** interface and can be used in Delphix Self-Service as data sources. You can use any data type supported by the Delphix Engine as a data source. For more information, refer to the Linking/Provisioning documentation for the standard Delphix Engine. The following is an example of the many kinds of data sources you can use to create a data template.



When adding data sources to the data template, it is important to consider whether there are any dependencies between them. For example, do data operations need to begin with a VDB (database) source before the same operation occurs on vFiles (application binary)? Or can data operations be performed in parallel with each of the data sources? The data source dependencies are by default synced together in parallel during any data operation, including starting the data container and its sources. When working with specific ordering constraints, such as with Oracle EBS, you can set up and configure the ordering sequence for each data source.

Procedure for Creating a Data Template

To create a data template:

1. From the drop-down menu in the upper right-hand corner of the Delphix Management application, select **Self-Service**.
2. On the **Overview** page, click **Add Template**.

ADD TEMPLATE +

This will send you to the **Create Data Template** page.

3. Enter a **Name** for the data template.
4. Optionally, enter a **description** for the data template.
5. Click **Add Data Source** to add data sources to the template.
Each data source name will include the name of the datasets group with which it is associated.

Create Data Template

Name
Shazam

Description

Order
 Set startup order of data sources

Data Sources

Name
Test

Replica
Default

dSource, VDB or vFiles
Boomerang (Avengers 2.0)

Notes

ADD DATA SOURCE + CANCEL CREATE

Create Data Template window with data source drop-down menu

To select a replica data source, first, select the name of the replica it belongs to. Then pick the replica masked VDB from the drop-down menu.

Create Data Template

Name
Shazam

Description

Order
 Set startup order of data sources

Data Sources

Name	Replica
My template	Default
	dSource, VDB or vFiles
	Bats (Avengers 2.0)
	Boomerang (Avengers 2.0)
	Dash (Avengers 2.0)
	Flash (Avengers 2.0)
	Flash 2.0 (Avengers 2.0)
	HR 1.0 (Avengers 2.0)
	HR 2.0 (Avengers 2.0)
	test123 (Avengers 2.0)

Notes

ADD DATA SOURCE + CANCEL CREATE

- To set a startup order, select the Set startup order of data source checkbox, then from the drop-down select the order.
- Select **Create**.

⚠ Default vs. Setting the Ordering of Data Sources to a Data Template

You have the option of setting the ordering of data sources to a data template. This option executes Delphix Self-Service operations sequentially on each data source (rather than in parallel), to ensure consistency among sources that need to be started/shut down in a particular order. You cannot change this setting after the data template has been created. **If you want the default behavior of running operations in parallel, do NOT select the box labeled "Set startup order of data sources".**

When your template has ordering constraints, as with Oracle EBS, you must set the startup order for each data source. Select the **Set startup order of data sources** box. The Delphix Engine will select the data source with order 1 as the first source started and the last one to be stopped. The data source with order 2 will be selected as the second source started, and this sequence will continue until the last data source is selected and ordered. Note that it is not possible to have operations performed in parallel on a subset of data sources and sequentially on a different subset of data sources.

Configure the Synchronization and Consistent Ordering of Data Sources

1. Select the **Set startup order of data sources** box.
2. Use the drop-down menu to select the **source** you want to include. The drop-down menu will display all dSources in the system and all VDBs and vFiles that are not already assigned to a Delphix Self-Service data container.
3. Enter a specific **name** for the data source.
4. Optionally, enter a **description** in the **Notes** section. Delphix Self-Service users see a copy of these notes in the data containers they own.
5. Click **Add Data Source** to continue to add and configure more data sources to the data template. You can remove data sources using the **Delete** icon.
6. By default, the **Order** of the sources will correspond to the order they are added. You can also edit the **Order** using the dropdown for each source.

7. Click **Create** to finish creating the data template.

Create Data Template

Name

Description

Order
 Set startup order of data sources

Data Sources

Name

Replica

dSource, VDB or vFiles

Notes

Order

Name

Replica

dSource, VDB or vFiles

Notes


Order

ADD DATA SOURCE +
CANCEL
CREATE

⚠ For Oracle EBS, the vFiles dbTechStack will have order 1, the Oracle database order 2, and the vFiles appsTier order 3. For more information about EBS, see the EBS documentation. Once you have created a template, you cannot change the set of data sources in it. Any VDBs or dSources being used as data sources in Delphix Self-Service will appear with a special badge in the Management application.

Viewing Data Templates

As the Delphix Admin user, you can view what sources have been included in a data template. You can distinguish the masked sources from the unmasked sources by referring to the corresponding data icons, as seen below.



Example of a template containing both masked and unmasked VDBs

Managing Data Templates

After you have created the data template, it will be visible from the **Overview** page under the **Data Templates** tab, which is the default tab.


Name	Replica	Containers	Last Updated
Shazam		0	Wed, Oct 4, 2017 10:25:22 AM
Template 1		1	Wed, Oct 4, 2017 9:57:53 AM
Template 2		2	Tue, Oct 3, 2017 11:24:53 AM

Data Template Details in the Overview page

Notes


- Each line corresponds to a data template and contains high-level information about that data template. For example, the number of child data containers is visible in the **Containers** column.

- You can search, sort, and filter the list of data templates, making it easy to manage a large number of data templates in Delphix Self-Service.



Editing a Data Template's Name

1. Click the **Edit** icon next to the data template name.
2. Enter the new **name**.
3. Click the **checkmark** icon to confirm changes.



Deleting a Data Template

1. Select the **data template** you want to delete.
2. Click the **Delete** button in the upper right-hand corner.



Data Containers

If there are any data containers provisioned from the data template, you must remove them before you can delete the data template. See instructions in [Data Container Activities](#).

Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Data Container Activities](#)


Understanding How to Manage Data Template Details

The Data Template Details Page

In the **Overview** page, under the **Templates**, click the data template's **name**. This will direct you to the templates details page. You can use this page to view and configure details of an individual data template. It consists of a number of tiles, described below.

Summary

Use this tile to get an overview of the data template and its child data containers.




Summary Details for Templates

Notes

- The graphs labeled **Operation Counts By Week** and **Container Data Age Distribution** give a sense of the amount of activity in the data template over time
- **Top Checkouts** shows at a glance which bookmarks have been used most frequently as part of a **Restore** or **Branch** operation
- **Top Users by Data Operations** shows at a glance which users are the most active

Containers


Use this tile to create, view, and delete child data containers from this data template.



Container Details

Sources


In this tile, you can view the data sources that this data template uses. Each data source has a user-visible name, a description, and a set of properties that consist of arbitrary key/value pairs. This information will be included in the data containers provisioned from this template.



Sources Details

Properties

Use this tile to edit the data template's properties. Properties are arbitrary key/value pairs associated with the data template. These values will be propagated to all data containers provisioned from this template. This provides a way for you to annotate data templates and data containers with whatever information is relevant to their use case.




Properties Details

Bookmarks

Use this tile to create and manage bookmarks on the data template. A bookmark represents a given point in time that is protected against retention. Bookmarks created on a data template are visible to all of the data containers

provisioned from it. For more details, refer to the Bookmarks section in the [Delphix Self-Service Data User Guide](#).



Bookmarks Details

Capacity

Use this tile to get information about the storage associated with the data template and its child containers.

Related Topics

- [Delphix Self-Service Admin Guide](#)

Understanding Data Containers

Delphix Self-Service Data Container Overview

Data containers are provisioned from data templates by administrators and assigned to a user. A data container represents a socket that is capable of making any data within the data template accessible. The user controls what data they want to access. Delphix Self-Service users have effectively been provisioned a set of "physical" resources, such as a database on a host that consumes some set of resources. A data container is comprised of a virtual database (VDB) or vFiles provisioned from each source in the data template from which it is created. The data container manages these VDBs, and the data operations performed on a data container will only impact these VDBs. Data containers represent the separation between IT infrastructure and end-users. IT determines the set of VDBs or vFiles to allocate to a data container, and users determine the data that they want accessible in the containers allocated to them.

Data containers can be used to access any data within a single data template, but not across templates. Users have the ability to populate the data within their data container from any point in time on the data template, the data container's history, or shared bookmarks from other data containers. Although operations are all accomplished by performing Timeflow operations on the underlying VDBs, the data containers hide the VDBs and their underlying properties from users. None of the data container operations require provisioning additional VDBs; everything is accomplished using the resources assigned when the data container is created.

Operation	Description
Refresh	This is the same basic concept as Refresh in VDBs. In Delphix Self-Service, Refresh will update the data on the active branch of a user's data container. The user will then have the latest data in the sources of the data template from which the container was provisioned.
Restore	Restore allows a Delphix Self-Service user to update the data on the active branch of their data container to one of the following: <ul style="list-style-type: none"> any point in time on the data container the data template from which the container was provisioned a bookmark This operation effectively means, "Take me to the data at this time."
Reset	Reset is a simplified version of Restore built to support the notion of "undo." It allows a user to reset the state of their application container to the latest operation. This can be useful for testing workflows where, after each test, users want to reset the state of their environment.
Branch	A branch represents a logical timeline, effectively a task on which a user is working. Only one branch can be active at a time, but a user can use multiple branches to track logically separate tasks. Delphix Self-Service branches do not require the allocation of a new VDB; instead, they are comprised of a collection of TimeFlows within a VDB.
Activate	This allows the user to select which branch they want to be active. Only a single branch within a data container can be active at a time.

Operation	Description
Bookmark	This creates a semantic name for a point in time and prevents this data from being removed by the retention policy. Bookmarks can be annotated with tags to make them easier to search for. In addition to tags, bookmarks allow a user to enter a description of what the bookmark represents.
Share	Bookmarks can be shared, which allows them to be seen by users who own data containers that have been provisioned from the same data template. This allows users to share data, providing a way for other users to either restore their existing timeline or create a new branch from these shared points.

Once a data operation has been selected its progress can be viewed from the Action sidebar in the Management application. From the Action sidebar, users cannot cancel the enable, disable or activate branch operations.


Automatic Retry for Data Operations

To make operations on data containers more robust, Delphix Self-Service supports automatically retrying failed sources during data operations. You can specify a maximum number of retry attempts so that if an operation fails on any individual data source within a data container, it will be automatically retried until it succeeds, or until the retry limit is reached.

Automatic retry applies to any Delphix Self-Service operation that changes the data in the data container, such as Refresh, Restore, Reset, or Create Branch. This setting can be especially useful in scenarios where there are a large number of sources in a data container, and some sources fail to update the first time. If the reason for the failures was intermittent, an automatic retry may allow the sources that failed to succeed, and the operation can still complete successfully. The default number of retry attempts is **1**.

To change the number of automatic retries:

1. Click the user icon in the upper right-hand corner of the screen.
2. From the drop-down menu, select **Settings**.
3. In the field **Automatically retry data sources...**, enter a new maximum; alternatively, use the arrows to increase or decrease the number.
4. Click **Save**.



Settings in the user drop-down menu

Delphix Self-Service Data Container Recovery


Data Containers Consistency

Delphix Self-Service allows you to group multiple datasets in the same data container. This makes it easy for you to access entire applications such as PeopleSoft, including binaries and code. If a data container represents an application, then there are likely to be dependencies between the application's datasets. For example, the vFiles data source containing the code will depend on a specific version of the database's schema. Therefore, it is important that all dataset sources are drawn from the same point in time. If they are, the data container is in a "consistent" state; if they are out of sync, or "inconsistent," errors will occur. For example, if the vFiles data source containing the code has been updated more recently than the database's schema, the dependency cannot work.

Delphix Self-Service currently has no way to determine whether the application is consistent. However, it attempts to minimize the chance that dataset sources are out of sync whenever it performs a data operation such as refresh, restore, or reset. When performing a data operation, Delphix Self-Service attempts to snapshot all dataset sources from a point in time as close as possible to the desired time. If at least one of the data sources fails to go to the desired point, then Delphix Self-Service considers the data container to be in an inconsistent state. The application as a whole may still be working, but Delphix Self-Service assumes that the failed dataset's data is not the correct version. To return to a consistent state, you must perform a recovery operation on the data container.

Data Container Recovery

Prior to performing any data operation, Delphix Self-Service takes snapshots of all datasets. Recovery is the process of rolling back a data container to a snapshot, thereby restoring it to a consistent state. When a failure occurs, you will see the following screen:



You can either perform recovery or use a different data container. Whether the recovery will fail or succeed depends on exactly why the data operation failed in the first place. If the problem was intermittent, such as a temporary network problem causing SSH failure, then performing recovery should work. If the problem is persistent – for example, the target host is out of space – then intervention is required; recovery will not succeed until you address the underlying root cause of the failure.

Admins can see the underlying failure in the **Actions** sidebar or the **Job History** dashboard. The **Actions** sidebar is the preferred place to view the failure; it has a hierarchical display that makes diagnosing the failure more straightforward.


Preserving Independent Containers in Delphix Self-Service During Replication

Replication is used for data backup and recovery as well as for managing and sharing data across remote data centers. Delphix Self-Service users can preserve their data after replication jobs. In the past, if replication occurred on templates in containers, users would lose the data in their containers. Now admin users can preserve containers to be used independently of replication jobs.

Independent containers behave in the same way as other containers, with two exceptions:

- You cannot refresh them.
- The bookmarks created on them cannot be shared, because they do not have a template reference.

The functional overview of independent containers seen below represents the flow of steps between the source engine and the target engine. A description of what is occurring between each of the steps appears below the diagram.



Functional Overview of Independent Containers

- 1 In Delphix Self-Service, you can create a template on the source engine and then replicate the template to the target engine.
- 2 In the target engine, an admin can use the replicated template to create new containers and assign them to users. You cannot change the replicated template's name or the names of the containers with which it was replicated over.
- 3 Due to an update, the replicated template is deleted from the source engine.
- 4 The deleted replicated template will be removed from the target engine. Any new container created in step 2 loses the reference to the deleted template and becomes an independent container.

Creating Independent Containers

Prerequisites

- The replication source and the replication target must be running identical versions of the Delphix Engine – for example, Delphix Engine version 5.1.
- The target Delphix Engine must be reachable from the source engine.
- The target Delphix Engine must have sufficient free storage to receive the replicated data.
- The user must have administrative privileges on the source and the target engines.

For more information, see [Configuring Replication](#) and [Understanding Data Templates](#).

Limitations of this Functionality

You can find independent containers in Delphix Self-Service on the target engine under the **Independent Containers** tab. They have the following characteristics:


- They cannot be refreshed, because they are no longer bound to a template.

- You can create bookmarks on them, but you cannot share those bookmarks because there is no common template.
- You can use them for branching, restoring, resetting, starting, and stopping.


Procedure

To create an independent container, complete the following steps:


1. On the source engine, create a template with a container.



2. From the user drop-down menu, select **Management**.



3. From the **System** menu, select **Replication**.



4. Next to **Replicated Profiles**, select the **plus** icon to **Create New Profile**.
5. Under **Objects Being Replicated**, select your Self-Service **template** and its associated **container**.
6. Enter your profile information.

7. Click **Create Profile**.


When replicating templates, you can select all, some, or none of their associated containers in the replication profile. This is done by selecting the checkbox next to the container's name in the **Create New Profile** window. When replicating a container, you must also replicate its associated template. Replicated objects cannot be modified on the target engine unless they are failed over, so you cannot modify the names of replicated data containers and templates.

8. Once the profile has been created, click **Replicate Now**.
9. On the target engine, click the user menu.
10. Select **Self-Service**.
11. The replicated template will appear in Self-Service on the target engine. The replica name is displayed next to the template name. You can edit regular templates by clicking the pencil icon next to the template name.


Name	Replica	Containers	Last Updated
Rep Template 1	nameOfReplica.dc1	0	Thu Oct 06 2016 13:04:30 GMT-0700 (PDT)
Rep Template 1		1	Thu Oct 06 2016 13:53:34 GMT-0700 (PDT)
Template 2		2	Wed Oct 05 2016 13:15:28 GMT-0700 (PDT)

12. Select the replicated template, then select **Containers**.
13. In the **Containers** window, click **Add Container**. In order to complete this action, you will need to ensure that there is data available from each data source in the template. This means that VDBs must have been provisioned from each replica dSource or VDB in the template. After the container is created, your replicated

template should have the new container you just created and the original container created in step 1.



14. On the source engine in Self-Service, delete your template.
15. From the user menu, select **Management**.
16. From the **System** menu, select **Replication**.
17. Replicate your profile to create a new independent container.
18. On your target engine, select **Self-Service**.
19. The new container is created. To find it:
 - a. Login to the target engine.
 - b. Click the user menu.
 - c. Select **Self-Service**.
 - d. In the Overview page, select the **Independent Containers** tab.



Related Topics

- [Configuring Replication](#)
- [Understanding Data Templates](#)
- [Delphix Self-Service Admin Guide](#)

Delphix Self-Service Data Container Activities

Configuring Data Containers in Delphix Self-Service

A data container is comprised of a set of virtual databases (VDBs), where each VDB is a direct child of the dSource, VDB, or vFiles in the data template's data sources. Delphix Self-Service does not automatically provision VDBs when creating a data container; a Delphix admin must create the required VDBs using the **Management** application.


i Info

Once a VDB or vFile is part of a Self-Service container, you cannot use the Management Service to rewind, refresh, or delete it. You can still use the Management Service to disable or enable it, take a snapshot of it, or provision a new VDB or vFile from it.

1. Select the **Overview** page.
2. Select a **template** from which you want to create the data container. This will take you to the **Data Template** page.

Add a Data Container

1. Click **Add Container** in the upper right-hand corner of the screen.



Details Panel and Dashboard

This will take you the **Create Data Container** page.

Create Data Container
 Template 2

Name

Owners

No Owners Added

Description

Data Sources

Refresh data sources to most recent template state
 Add data sources to container as-is

<p>Data Source Name</p> <p>dbdhcp3-bbdhcp-7RSG-1298.dc2.delphix.com dSource dbdhcp3-bbdhcp-7RSG-1298.dc2.delphix.com</p>	<p>VDB or vFiles</p> <div style="border: 1px solid #ccc; padding: 2px;">Vdbd_A6C (Untitled) ▼</div>
<p>Connection Info</p> <p>Host bbdhcp-7RSG-1298.dc2.delphix.com DB Name Vdbd_A6C DB Version Oracle 11.2.0.2.0 Oracle Home /u03/app/ora11202/product/11.2.0/dbhome_1 JDBC jdbc:oracle:thin:@10.43.75.172:1521/Vdbd_A6C</p>	<p>Description</p> <div style="border: 1px solid #ccc; height: 40px; width: 100%;"></div>

CANCEL
CREATE

Data Container Dialog


2. Optional: Enter information about the data container, such as the **Name** and **Description**.
3. Once the data container has been created Select the **Owners** for the data container from the search box. It is acceptable to have multiple owners for each data container. Any Delphix administrator is able to manage all containers, so the owners should be end-users. For details, see [Understanding Delphix Self-Service User Management](#).

Owners

Batgirl
delphix_admin

4. When a data container is created, you now have the option to:

- a. **Refresh data sources to most recent template state** – This option will refresh VDBs before adding them to the container. This is done to enforce that when multiple sources are used in a container, the sources are consistent.
 - b. **Add data sources to container as-is** – This option will not refresh the data sources.
5. Select the **VDBs** to use for this container's data sources. The available VDBs have the following constraints:
 - They have been provisioned from the dSources/VDBs belonging to the parent data template
 - They are not already part of another data template or containerIf there are no VDBs that meet these constraints, you may see a message informing you that you do not have any compatible VDBs.



VDB Warning Alert

6. Click **Create**.

Selecting Masked Data Sources for Data Containers

Prerequisites

- [Using Masked Data Sources with Delphix Self-Service](#)
- [Selecting Masked Data Sources in Data Templates](#)

Procedure

Once you have selected a child masked VDB for the data container, you can see the parent-child relationship as a masked source under data sources.

Data Sources

Refresh data sources to most recent template state (ensures data consistency)
 Add data sources to container as-is (possible that data between sources may be inconsistent)

Data Source Name	VDB	VDB or vFiles
Masked	dbdhcp2-dbdh_FWO-1489717771234	Vdbd_065 (Untitled)

Parent Masked VDB (points to dbdhcp2-dbdh_FWO-1489717771234)

Child Masked VDB (points to Vdbd_065 (Untitled))

Connection Info

Host kgbdhcp-tgt.dc2
 DB Name Vdbd_065
 DB Version Oracle 11.1.0.7.0
 Oracle Home /u02/app/ora11107/product/11.1.0/db_1
 JDBC jdbc:oracle:thin:@10.43.76.8:1521/Vdbd_065

Description

Order
Parallel

Data Source Name	VDB	VDB or vFiles
Unmasked	dbdhcp1-dbdh_IEG-1489715998270	Vdbdhcp1dbdh_IEG1489715998_38E (Untitled)

Connection Info


Host kgbdhcp-tgt.dc2
 DB Name Vdbdhcp1dbdh_IEG1489715998_38E
 DB Version Oracle 10.2.0.5.0
 Oracle Home /u01/app/ora10205/product/10.2.0/db_1
 JDBC jdbc:oracle:thin:@10.43.76.8:1521/Vdbdhcp1dbdh_IEG1489715998_38E

Description

Masked Data Sources Parent/Child Relationship

As an admin user, you can select both masked and unmasked data sources in both Delphix Self-Service templates and data containers.

Delphix Self-Service users will not know whether the data in their containers and branches is masked or unmasked. All Delphix Self-Service functionality remains the same regardless of whether a data source is masked or unmasked.



The figure above is an example of a data container with masked data.

Delete a Data Container

All data sources (VDBs and vFiles) in a Data Container are not deleted as part of the Data Container deletion process.

When performing the **Delete Container** operation, you can check the **Delete associated VDBs and vFiles** box in the dialog window to delete the data sources associated with the container.

Data Management Operations

Start a Data Container

Starting a Data Container does the following:

- Starts the data sources, This means that each data source listed in the **Source Details** section of the **Data Container** page will start using CPU and network resources on the host system it is running.
- Puts a copy of the data from the active branch into those data sources.

On the **Self-Service Toolbar**, click **Start**.

Stop a Data Container

Stopping a data container does the following:

- If not already done, copies the current data in the data sources into the active branch of the data container
- Shuts down the data sources. This means each data source listed in the **Source Details** section of the **Data Container** page will stop using CPU and network resources on the host system.

On the **Self-Service Toolbar**, click **Stop**.

Other operations on the data container, such as Stop, Reset, and Refresh, must be performed from the **Data Management** page:

Data Management Interface Shortcut in Delphix Self-Service Data Template

Locking a Container

On the Self-Service Toolbar, click **Lock**.




Locking a data container does the following:

- You become the only user who can perform operations on it.
- For all other users, the container appears disabled.

Unlocking a Container

Unlocking a data container does the following allows other Self-Service users to perform operations on that container.

This operation is only enabled if a container is currently locked. Only the user who locked a container or a Delphix Self-Service administrator can unlock it.



Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Understanding Data Containers](#)
- [Using Masked Data Sources with Delphix Self-Service](#)
- [Understanding Delphix Self-Service User Management](#)


Using Masked Data Sources with Delphix Self-Service

SDD Overview

You can now replicate masked data in a VDB directly to a target Delphix Engine without transmitting the unmasked data in its parent source. This is called [selective data distribution](#) (SDD). Although you can run selective data distribution ad hoc, it is typically run according to a predefined schedule. In the current release, there are some best practices and limitations you should know about before using data from SDD in Delphix Self-Service. This section covers the workflow for using SDD replicated data in templates and containers. It is aimed at administrators who are familiar with the process of creating a masked VDB, SDD replication and setting up objects.

Configuring Delphix Self-Service with Masked Data Sources

The diagram below illustrates the steps for using masked data sources in Delphix Self-Service.



Step 1: Provision a masked VDB on the source.

Step 2: Use SDD to replicate masked data to the target.

Use SDD to replicate your masked VDB to the target. The target VDB will be called the “replica masked VDB.”



Replica Masked VDBs

To keep the replica masked VDB up to date, configure a refresh policy for how often it should refresh. The refresh policy should be related to the schedule for SDD updates from the source. Refreshing more frequently will result in the VDB being unavailable to Delphix Self-Service more often than needed.

Step 3: Create a data template on the target.

To create a data template:

1. From the drop-down menu in the upper right-hand corner of the Delphix Management application, select **Self-Service**.
2. On the **Overview** page, click **Add Template**.

ADD TEMPLATE +

Add Data Template

This will send you to the **Create Data Template** page.

3. Enter a **Name** for the data template.
4. Optionally, enter a **description** for the data template.
5. Click **Add Data Source** to add data sources to the template.
Each data source name will include the name of the datasets group with which it is associated.

Create Data Template

Name
Shazam

Description

Order
 Set startup order of data sources

Data Sources

Name
Test

Replica
Default

dSource, VDB or vFiles
Boomerang (Avengers 2.0)

Notes

ADD DATA SOURCE + CANCEL CREATE

Create Data Template window with data source drop-down menu

To select a replica data source, first, select the name of the replica it belongs to. Then pick the replica masked VDB from the drop-down menu.



6. To set a startup order, select the Set startup order of data source checkbox, then from the drop-down select the order.
7. Select **Create**.

Step 4: Provision child VDBs from the replica masked VDB.


1. Login to the **Delphix Management** application for the target host.
2. Click **Manage**.
3. Select **Datasets**.
4. Select the **replica** that contains the dSource or VDB to be provisioned.
5. The provisioning process is now identical to the process for provisioning standard objects.

Step 5: Add data containers and select the child VDBs as data sources.

Follow the instructions to [add a data container](#).

Select a masked child VDB as a source for the container. As an admin user, you can select both masked and unmasked data sources in templates and data containers.

Once you select a child masked VDB for the data container, you can see the parent-child relationship as a masked source under data sources.



Masked Data Sources Parent/Child Relationship

Refreshing Masked VDBs in Delphix Self-Service Data Templates

Make sure that in step 4 above, you select the replica masked VDB as the source for the data template.



In order for new data to be available in the template on the target, you must do the following:

1. Refresh the masked VDB on the source. This will re-run the masking job.
2. After the refresh completes, execute the SDD spec for the masked VDB.

⚠ Wait to refresh
 If you do not wait until the refresh is complete, unmasked data may be sent to the target.

 **SDD update**

Although you can employ a policy to drive refreshes of the masked VDB, you cannot use that policy to drive the SDD update as well. You may need a combination of policies and scripts to automate the workflow.

Limitations

 **SDD Replication Profile**

You cannot add data templates to an SDD replication profile. As a result, you must create the data template on the target. This is step 3 above.

Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Understanding Data Containers](#)
- [Delphix Self-Service Data Container Activities](#)

Understanding Delphix Self-Service User Management

User Management Activities

This document describes the process of creating a user and assigning that user to a data container. It also provides an overview of the **User Details** page.

Creating a User

Follow the same process when creating a new user or modifying an existing Delphix user. Delphix Self-Service users do not have access to the existing admin user interface, and they can only access the **Data Container** page for containers they own.

1. From the **Management** application, select **Manage**.
2. Select **Users**.
3. Click **+** to add a new user or to make an existing Delphix Engine user a Delphix Self-Service user, select the user from the list.
4. Enter the appropriate information.
5. From the **User Type** drop-down menu, select **Self-Service Only**.



6. Press **Submit**.

The user is now a Delphix Self-Service user! They can now login to the Delphix Self-Service user interface, and you can make them the owner of a data container.

Notes

- Users will only be able to access the **Data Management** page. They will not be able to access the other portions of the Delphix Self-Service interface, nor the **Management** application.
- A Delphix admin user cannot be made a Self-Service Only user. However, admins can still use Delphix Self-Service and own a data container. Admins are also able to manage all data containers.
- A user who owns one or more data containers cannot be deleted.
 - For the list of data containers that a given user owns, see **User Details**.
- You cannot revoke a user's Self-Service Only role if they own any data containers.
 - For the list of data containers that a given user owns, see **User Details**.

Assigning a User to a Data Container

This section describes how to assign a user (created in the previous section) to a data container. Making a user the owner of a data container allows them to perform operations such as **Refresh** on that data container. Users cannot see or manipulate data containers that they do not own. You can either assign a user when creating a new data container or modify the owner of an existing data container.


Case 1: Assigning a user to a new Data Container

1. Navigate to the **Overview** page.


2. Select a **template** from which you want to create the data container.
This will take you to the **Data Template** page.
3. In the upper right-hand corner of the screen, click **Add Container**.
4. Optional: Enter information about the data container, such as the **Name** and **Description**.
5. Select the **Owners** for the data container from the search box. It is acceptable to have multiple owners for each data container.
6. Select the **VDBs** to use for this container's data sources and Save.

Case 2: Changing the owner of an existing Data Container

1. On the **Management Overview** page, select the **data template** from which the data container was provisioned.
2. Click the **Containers** tile in the left-hand panel.
3. Select the **trash can** icon next to the username.



4. Click the **Plus** icon.



5. Select the desired **owner** from the drop-down list.

The user you selected is now the owner of the data container and can perform operations on that data container.

User Details Page

This section provides an overview of **Delphix Self-Service User Details**. This page displays graphs related to the user's Delphix Self-Service activity, as well as a list of all of the data containers that the user owns.


1. On the **Management Overview** page, click the **Users** tab.



2. Select the **name** of the desired user to go to their **User Details** page.

The following sections will display important information:

Section	Information displayed
Operation Counts By Week graph	Shows the aggregate of all Delphix Self-Service operations performed this user has performed on all of their containers.
Container Age Distribution graph	Shows the average time since a data operation was performed on all of the user's containers
Containers	Lists all containers that the user owns



User Details

Related Topics

- [Delphix Self-Service Admin Guide](#)

Working with Multiple Container Owners

Delphix Self-Service administrators can designate multiple users as owners of a single data container. These users all share access to the same data container which means actions taken by one user will impact all users on the same data container. For example, if User A activates Branch X, User B will also see Branch X as the active branch. This ability for one user's actions to impact another user on the same containers creates new concerns for users sharing the same container. As a result, more processes should be put into place in order to coordinate usage between users. Each team is different, but strategies include:

- designating a person to perform certain data operations
- saving your work with a bookmark or creating/working on a personal branch
- being aware of who is using your data container/data before performing operations
- locking a container to prevent others from performing any operations on it

How many owners should a container ideally be shared between?

There is no technical limit built into the software, but it is best if a team of 5-10 users shares a single data container. In most cases, having fewer owners minimizes overhead and conflicting usage. One owner per container provides maximum productivity and minimal overhead, so this feature should only be used if your infrastructure or processes require that multiple users share a container. Additionally, Jet Stream Only users currently cannot see other users with whom they share the container.

How should users handle potentially disruptive operations?

If one user performs an operation on a data container, it will affect the other owners of that container. Additionally, each user has permission to perform the same operations on the data container; currently there are no fine-granularity permissions that limit the operations a user can perform. All operations are potentially disruptive, but the level of disruption varies by operation. If any of the following operations are performed at the same time, the second operation will fail due to a conflict when processing the job.

Conflicting operations

- Refresh
- Restore
- Reset
- Enable/Disable
- Create Branch
- Activate Branch
- Delete Branch
- Create Bookmark
- Delete Bookmark

If User A performs a destructive operation while User B is "using" the data container, the operation will destroy User B's current state. Currently, the interface does not provide insight into whether the data container is in use by another user.

Destructive operations

- Refresh
- Restore
- Reset
- Enable/Disable
- Create Branch
- Activate Branch

Deleting objects

All owners can delete any bookmarks or branches in the container, regardless of who created them.

Coordinating Users

Opportunity for disruption increases as more owners are sharing a single container. Sharing a container works best when users can communicate with each other, such as when they are part of a team, or when they are working with the container at different times.

Additionally, these disruptions can be avoided when Delphix Self-Service users lock their containers to prevent other users from performing operations on them. Users cannot see the other users with whom they share the container. However, if a user locks a container, only that user will be able to use the container; it will appear disabled to others.

What operations could disrupt others using a container?

Potentially disruptive operations include:

- Refresh
- Switching active branches
- Deleting bookmarks
- Creating Branches
- Un-sharing bookmarks
- Restore
- Reset
- Starting/ stopping your container

What processes should I put in place?

The more owners you assign to a single container, the more processes you should put in place to coordinate usage between users. Each team is different, but strategies include:

- designating one person to perform certain data operations
- saving your work with a bookmark or a creating and working on a personal branch
- being aware of who is using your data container before performing operations
- locking a container to prevent others from performing any operations on it

Where can I see which user has performed what operation?

You can see which user has performed which action in the **History** tab of the **Data Operations** screen.

Be aware that operation counts in the template view are currently tabulated based on the container, not the user performing the operation.

DATA OPERATIONS ▸ Template 1 Container 1 ★ default

Filter

Action	Description	Time
Refresh	User "delphix_admin" refreshed Jet Stream data container "Container 1".	Wed, Oct 4, 2017 12:53:44 PM
Create Branch	User "delphix_admin" created the new branch "default" on Jet Stream data layout "Container 1".	Tue, Oct 3, 2017 11:23:07 AM

TIMELINE
SUMMARY
SOURCES
HISTORY
BOOKMARKS

Data Operations history tab

Where can I see which containers are unlocked/locked?

Unlocked containers can be viewed by typing “unlocked” in the search filter.


Template 2

Containers ADD CONTAINER + Filter

Name	Operations	Last Updated	
Container 2	2	Fri, Jul 28, 2017 14:57:26 PM	Container 3
Container 3	1	Thu, Jul 27, 2017 18:16:25 PM	Owners + No Owners default Data Management

SUMMARY
CONTAINERS
SOURCES
PROPERTIES
BOOKMARKS

You can view locked containers by typing “locked” in the search filter. To find containers locked by a specific user, type “locked by {username}” in the search filter.



Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Delphix Self-Service Data Container Activities](#)
- [Understanding Data Containers](#)

Understanding Bookmarks

Bookmarks Overview

Bookmarks are a way to mark and name a particular moment of data on a timeline. You can restore the active branch's timeline to the moment of data marked with a bookmark. You can also share bookmarks with other users, which allows them to restore their own active branches to the moment of data in your container. The data represented by a bookmark is protected and will not be deleted until the bookmark is deleted. To help manage the space used by this data, users can set an optional expiration date for a bookmark. At the end of the set date, the bookmark will automatically be deleted. Once created, you can easily locate a bookmark through one of the bookmark viewers in the interface. To understand how to use bookmarks, please refer to the [Delphix Self-Service Data User Guide](#).

Using Bookmarks in Data Templates

An admin user can create a bookmark on a template that will then be automatically shared to all containers created from that template. Additionally, an admin user can create a bookmark on the master template timeline with the point of time you are interested in. The bookmark will always be saved from retention policies, and a new branch can be created from this bookmark.

Default Bookmark Expiration

You can set a value that controls the default expiration time, in days, for Bookmarks. This setting only applies to new bookmarks that are created through the Delphix Self-Service application, not the CLI or API. Note that this only controls the default selection; users can still disable expiration or pick a different date for a specific bookmark if they wish. This setting is disabled by default.

The screenshot shows the Delphix Self-Service interface. At the top right, there is a user profile for 'admin' and a navigation menu with options: Overview, Data Operations, Usage, Settings (highlighted), Management, Engine Setup, and Logout. Below the navigation is a search bar and buttons for 'DELETE' and 'ADD TEM'. The main area contains a table with columns: 'e', 'Replica', 'Containers', and 'Last Updated'. The table has two rows of data, both with red warning icons in the 'Replica' column. The first row shows '1' container and 'Tue, Nov 6, 2018 9:58'. The second row shows '2' containers and 'Tue, Nov 6, 2018 10:04:46 AM'. A 'Settings' modal is open in the bottom left, showing two checked options: 'Default expiration for bookmarks: 30 days' and 'Automatically retry data sources that fail during Self-Service data operations: 1 retries per source'. The modal has 'CANCEL' and 'SAVE' buttons.

e	Replica	Containers	Last Updated
	⚠	1	Tue, Nov 6, 2018 9:58
	⚠	2	Tue, Nov 6, 2018 10:04:46 AM

Related Topics

- [Delphix Self-Service Admin Guide](#)
- [Delphix Self-Service Data User Guide](#)

Understanding Delphix Self-Service Usage Management

Usage Management Dashboard Overview

Data templates are comprised of dSources, virtual databases (VDBs), and vFiles. These data sources are controlled by the standard policies configured in the **Management** application of the Delphix Engine. As with existing containers, space will be reclaimed by the retention policy over time. As retention cleans up historical data, users will no longer be able to use those points in time to restore or branch. In Delphix Self-Service, an admin can create a bookmark on the data template timeline, which will prevent retention from cleaning up the data that a bookmark references.

Data containers are comprised of VDBs provisioned from the sources defined in the data template. Similar to VDBs in the **Management** application, data containers' VDBs will share blocks with the source from which they are provisioned. This prevents the referenced data on the source from being cleaned up by retention. Retention for these VDBs is controlled by the standard Delphix Engine retention policies. As on templates, bookmarks in data containers will prevent storage from being reclaimed by retention. In addition, Self-Service will ensure that the latest data on each branch is never removed.

The **Usage** pages of the data templates and data containers provide information that can help you understand how storage is being used, how to reclaim space, and how much space you are able to reclaim.


Usage Overview is a top-level page, along with the **Data Operations** and **Overview** pages. It contains the space usage breakdowns by data templates and users.



Template Usage Overview

The **Template Usage** page, seen in the image below, contains the usage breakdowns for data templates and users. The interface is interactive and allows you to visualize data by interacting with pie charts, bar graphs, and tables. The pie chart contains information about the top 10 space consumers; the table at the bottom contains information about all of the templates and/or users.

The table below the charts includes category fields. You can find corresponding descriptions by hovering over the names of the fields in the table:



The Template Usage page

Additionally, the table allows you to sort, navigate, and interact by clicking the field category of interest. For example, to sort the table, click a **column header** such as **Unvirtualized** and the table will sort by that category. To navigate to a particular data template or user, you can click either the **pie slice** or the **name** of the template/user in the table.

A close-up view of the 'All Templates' table. A tooltip is displayed over the 'Containers' column header, explaining that it represents the amount of space used by data containers provisioned from the template, which will be freed up if deleted or purged.

Template	Total	Bookmarks	Containers	Unvirtualized
Template 1	1.81GB	127.1MB	1.6GB	17.07GB
Template 2	1.19GB	0B	1.1GB	
Testing	0B	0B		

Table of templates/users


The field categories display the following information:

- **Total** – The sum of the space used by the data containers provisioned from this data template and by the bookmarks created on this template. This is the space that will be freed if you delete the template.
- **Containers** – The amount of space used by the data containers provisioned from this data template. This is the space that will be freed if you delete or purge all of the data containers.
- **Bookmarks** – The amount of space used by the bookmarks on this data template. This is the space that will be freed if you delete all bookmarks on the template.
- **Unvirtualized** – The amount of space that would be used by the data in this template and its child data containers without Delphix virtualization.

The pie chart and table graphs can help you analyze storage usage information.

Template Usage Details


The **Usage** tile appears at the bottom of the Self-Service navigation sidebar, as seen in the image below. Usage summaries are available for templates, containers, and users. For example, when you click the **Usage** tile on the **Template Details** page, the usage details you interact with will be in the context of the selected data template.



The Usage tile

User Usage Overview

The **User Usage Overview** page provides graphical visualizations of space used by users assigned to data containers.



The field categories display the following information:

- **Referenced** – The amount of space used by data containers that are owned by this user. This excludes the space that this user is sharing with other users.
- **Containers Owned** – The number of data containers owned by this user

Template Usage (Containers) Overview

The **Template Usage Details** page, as seen below, shows the space used by data containers provisioned from the template and the bookmarks created on the template.



Container Usage

The stacked bar graph shows information about the top 10 space users. You can re-sort the graph based on the fields in the **Sort by** the legend on the top right-hand corner of the screen as seen in the image above. For example, if you want to know which data containers are sharing the most data with others, you can un-select **Shared (others data)** and **Unique** by clicking them in the legend.



The Sort by legend

i Legend Items


When the legend items are not selected, their corresponding colored boxes turn gray and the data is removed from the chart. The data and name will reappear when you re-select by clicking the grayed-out category you want.

The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this data container. This assumes that it also deletes underlying data sources.
- **Shared (others data)** – The amount of space that cannot be freed on the parent data template (or sibling data containers) because it is also being referenced by this data container due to Restore or Create Branch operations. The snapshots on the template or sibling container are what use up space.
- **Shared (self data)** – The amount of space that cannot be freed on this data container because it is also being referenced by sibling data containers due to Restore or Create Branch operations, via shared bookmarks
- **Unvirtualized** – The amount of space that would be used by the data in this container without Delphix virtualization

Template Usage (Bookmarks) Overview

As shown in the image below, the **Template Usage Details** page provides the usage information about bookmarks created on a template. The primary categories of information include **Unique**, **Shared (others data)** and **Shared (self data)**.




Template Usage (Bookmarks)

The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this bookmark
- **Shared** – The amount of space referenced by this bookmark that cannot be freed by deleting this bookmark because it is also referenced by neighboring bookmarks or branches that have been created or restored from this bookmark
- **Externally Referenced** – The amount of space referenced by this bookmark that cannot be freed by deleting this bookmark because it is also being referenced outside of Self-Service – for example, by a retention policy.

Container Usage (Branches) Overview

The **Container Usage Details** page shows the usage information about the branches and bookmarks created on a container. The primary categories of information include **Unique**, **Shared (others data)**, and **Shared (self data)**.



The Container Usage Details page

The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this branch
- **Shared (others data)** – The amount of space that cannot be freed on the parent data template or sibling branches because it is also being referenced by this branch due to Restore or Create Branch operations. The snapshots on the template or sibling container are what use up space.
- **Shared (self data)** – The amount of space that cannot be freed on this branch because it is also being referenced by sibling data containers due to Restore or Create Branch operations, via shared bookmarks.

Related Topics

- [Delphix Self-Service Admin Guide](#)

Delphix Self-Service Data User Guide

This section covers the following topics:

- [Welcome to Delphix Self-Service](#)
- [Delphix Self-Service Data Concepts](#)
- [Delphix Self-Service User Interface](#)
- [Understanding Timelines and How to Preserve Data in a Point in Time](#)
- [Data Container Activities](#)
- [Containers with Multiple Owners](#)
- [Working with Bookmarks in a Data Container](#)
- [Understanding Delphix Self-Service Usage](#)

Welcome to Delphix Self-Service

Delphix Self-Service grants access to the data that users need, whenever they need it. Once users have been assigned a data container, they can control the data available within it. This means they can refresh to the latest production data, roll back to a previous point in the data container's timeline, and share data with another user without requiring any involvement from Information Technology or database administrators (DBAs). Delphix Self-service data management allows developers to be more productive while using fewer resources, dramatically improving operational efficiency.

User Roles and Permissions

Self-Service has two types of users:

Admin User

Admin users have full access to all report data and can configure Delphix Self-Service. Additionally, they can use the Delphix Engine to add/delete reports, add/delete users, change tunable settings, add/delete tags, and create and assign data templates and containers.

Data User

Data users have access to production data provided in a data container. The data container provides these users with a playground in which to work with data using the self-service toolbar.

Login

1. Access Delphix Self-Service by opening a web browser using the **IP address** or **DNS qualified hostname**.
2. Login with the **User ID** and **Password** the Delphix Administrator has provided for you.

Changing Your Default Locale

As a user, you can change your default locale by doing the following:

1. Click the **user login** icon in the upper right-hand corner of the screen.
2. Click the **Locale** drop-down menu.
3. Select the desired locale.

Add User

User Details

User Privileges

User Details

Authentication Type

Username

Password

Confirm Password

Email Address

User Type

Locale

Last Name

User Profile window

Delphix Self-Service Data Concepts

Understanding Data Sources

A data source can be a database, an application, or a set of unstructured files. Delphix administrators configure the Delphix Engine to link to data sources, which pulls in the data of these sources. The Delphix Engine will periodically pull in new changes to the data, based on a specific policy. This, in turn, begins building a custom timeline for each data source. Additionally, the Delphix Engine can rapidly provision new data sources that are space-efficient copies, allowing users to work in parallel without impacting each other.

Understanding Data Templates

Data templates are the backbone of data containers. They are created by the Delphix administrator and consist of the data sources that users need in order to manage their data playground and their testing and/or development environments. Data templates serve as the parent for a set of data containers that the administrator assigns to users. Additionally, data templates enforce the boundaries for how data is shared. Data can only be shared directly with other users whose containers were created from the same parent data template.

Understanding Data Containers

A data container allows data users to access and manage their data in powerful ways. Their data can consist of application binaries, supporting information, and even the entire database(s) that underlie it.

A data container allows users to:


- Undo any changes to their application data in seconds or minutes
- Have immediate access to any version of their data over the course of their project
- Share their data with other people on their team, without needing to relinquish control of their own container
- Refresh their data from production data without waiting for an overworked DBA

A data container consists of one or more data sources, such as databases, application binaries, or other application data. The user controls the data made available by these data sources. Just like data sources in a template, changes that the user makes will be tracked, providing the user with their own data history.

The **Data Container Interface** lets users view the details and status of their data container and its associated data sources, as well as manipulating which data is in those sources. The **Data Container Interface** includes a section called the **Data Container Report Panel**, which displays details about each source, including the connection information needed to access it – for example, the java database connectivity (JDBC) string for a database. This connection of information is persistent and stable for the life of the data container, regardless of what data the resources are hosting.

Data Flow

The Delphix Self-Service data flow diagram below demonstrates how a Delphix Self-Service data user accesses data sources. Data sources are connected to a Delphix Engine, which is controlled by the Delphix administrator. The Delphix administrator will connect all data sources that developers and quality assurance (QA) teams need to a data template. This data template acts as a parent source to create the data containers that the administrator will assign to data users. Data sources flow from the Delphix Engine into a data template and downstream into a data container, where a data user or users will use the data sources to complete tasks. The data container acts as a self-contained testing environment and playground for the data user. Additionally, data users are able to set, bookmark, and share data points in their container with other data users of other data containers, as long as all the data containers were created from the same parent data template.




Data Flow

Understanding Branches

You can organize data in the data container into task-specific groupings, called "branches." For example, you can use a branch to group all the data you have used while addressing a particular bug, testing a new feature in an application, or exploring a business analytics scenario. By default, Delphix Self-Service automatically creates the first branch of source data for you when you login for the first time. You can view the default branch and any additional branches that you create over time by clicking the Branch tab. Additionally, to the right of the default branch, you will see an interconnected branch timeline unique to whichever branch is currently active. The illustration below displays both the default branch in the Branch tab of the Data Container View Panel and the default branch timeline.



Branch Tab




Branch View Panel and Branch Timeline

A branch is used to track a logical task and contains a timeline of the historical data for that task. One branch is the "active" branch, which means that it is the branch that is currently being updated with new data from the data sources. At any time, you can change which branch is active and thus change which data is in the associated data sources.

Delphix Self-Service User Interface

The User Interface is organized within a single web browser page. The screen serves as a data container report and management panel.

Data Container Report Panel



Data Container Workspace

Data Container Report Panel

The Data Container Report Panel for users consists of two tile buttons. They are summarized below as **Timeline** and **Bookmarks**.

Timeline

1

The **Timeline** tile allows you to view the timeline associated with a branch. Note that this only shows the timeline for a single branch. The branch timeline is how a user interacts with data in the container to mark, stamp, and perform tasks that occur at various points in time.


Bookmarks

2

The **Bookmarks** tile allows you to view and edit details about bookmarks within this data container and bookmarks accessible from it.

Data Container Workspace

The Data Container workspace is reached by selecting a template from the Management Overview page. Select a container and then select the Data Management link.



Data Container Workspace (Top Half)

Data Container Workspace

The Data Container Workspace contains all the tools, actions, and view panels needed to begin using Self-Service features. For example, the workspace allows a user to view the history of their data on a branch, and to refresh, reset, and restore that data.

User Login and Settings Drop Down Menu

7

The **user login** icon in the upper right-hand corner of the screen provides a drop-down menu with options to change your password and/or log out.

Data Container View Panel

8

The **Data Container View** Panel, found on the left-hand side of the screen, is divided into three tabular sections: time, branches, and bookmarks. These tabs allow you to find and select data that you are interested in. Based on user selections made in the view panel, the corresponding branch timeline can change.

Switch Container

9

Allows you to switch between containers.

Data Container Self-Service Toolbar

10

The Data Container Self-Service Toolbar allows you to perform tasks and activities with data in the current container, by clicking on the following user action icons:

- **Activate** will make a branch active
- **Bookmark** will mark an interesting point of data on a branch timeline
- **Branch** will create a branch that supports one task. A branch is a group of data time segments called a "timeline."
- **Share** will share a bookmark with users of other data containers from the same template
- **Refresh** will refresh each source in the data container on a branch timeline to the latest data in the corresponding source of the data template.
- **Restore** will restore the data to a point in time from the template, the container, or a shared bookmark.
- **Reset** will reset to the last interesting moment of data time on the current data timeline
- **Stop** will stop a data container
- **Start** will start a data container
- **Lock** will lock a data container for the current user
- **Unlock** will unlock a data container.

Branch Timeline

11

Use this to view the timeline associated with a branch. Note that this only shows the timeline for a single branch. The branch timeline is how a user interacts with data in the container to mark, stamp, and perform tasks that occur at various points in time.

Understanding Timelines and How to Preserve Data in a Point in Time

Understanding Timelines

Branch Timeline


A branch timeline acts as a dynamic point-in-time interface for user actions within the branch. You can interact with the source data in the active branch by using both the branch timeline and icons along the **Self-Service Toolbar** at specific points in time. Common activities include re-setting data sources to run a test, refreshing the data container with the most current source data, and bookmarking data to share or track interesting moments of time along the branch timeline. Users work with one branch at a time to perform a series of actions related to a particular testing or debugging task such as data updates or starting and stopping data. The initial branch on your data container may include data from before the data sources were made into a container. As you work within your data container, you can create more branches overtime to run or complete separate tasks. Additionally, the data container tracks each branch and the corresponding actions you perform on the branches. To view the actions completed over the life of a branch, see the container timeline in the **Time** tab of the **Data Container View Panel**.



Branch with Timeline Segments Over the Life of the Branch

Container Timeline

The **Time** tab displays the data container's timeline, which acts as a wall clock of time. It shows continuous real-time across all branches and timeline segments. You can scroll up and down in the container timeline to find the point in time that interests you.



Time Tab Timeline

Clicking a point in time in the container timeline will display the corresponding branch timeline capturing any actions performed on the branch. Additionally, should you need to select a time between tick-marks, you can use the **time input field** in the time selector on the left side of the screen.



Container Timeline

Selecting a Point in Time with the Time Selector

1. In the **time selector**, type in a **date** and **time** with the following format: Month/Day/Year Hour:Minute:Second{am|pm}. For example: 1/26/2015 1:14:13pm.
2. Press **Enter**.


The time input field will show the selected time. Now that you have entered the specific time you want, you can use the toolbar to select the data operation that you want to be performed at this point in time. Data operations can include Create Bookmark, Create Branch, and Restore.

Invalid Time Value

If you type in an invalid time value, or a time that is out of range, the value you typed in will revert to the previous default that existed before.

Selecting a Point in Time with the Time Selector Calendar

1. In the time selector, click the **calendar** icon to the left of the input field.
2. From the window that appears, click the **date** you want to use.
3. Select and **time** you want to use.




4. In the toolbar, click the button for the data operation that you want to perform at this point in time. Data operations can include Reset, Create Branch, and Create a Bookmark.

Picking a date

The flyout will not let you pick a date that is before the first point of data time in the container, or after the present moment.


Understanding the Self-Service Toolbar

The **Self-Service Toolbar** contains self-service action icons that represent available actions a data user can perform. You can distinguish between available and unavailable icon actions by the use of color on the toolbar. Actions available to you will be red, and actions that are unavailable will be grey. All actions are dynamic, and availability will change based on how you use and work with data in both the branches and data container(s) that are assigned to you.



Self-Service Toolbar

For example, your options for actions on the **Self-Service Toolbar** can change if the branch of the branch timeline you are working with is activated. In the illustration below, the screen shows a user working in an active branch. Notice the bright red star at the end of the timeline. This indicates that the branch is active. Also, notice which actions are and are not available to the user on the **Self-Service Toolbar**.



Self-Service Toolbar with a Point in Time selected on an Active Branch Timeline

The **Self-Service Toolbar** is dynamic and will change based on tasks a user performs in Self-Service. These workflows will influence how and when self-service actions become available on the Self-Service Toolbar.


Branch Timeline Segments

A branch timeline with segments is a visual representation of actions taken on a branch timeline over a time span. The timeline segments represents data in time that is no longer contiguous once a user clicks **Create Branch, Refresh, Reset, or Restore** on the **Self-Service Toolbar**. A vertical bar between each of the segments appears to remind a Delphix Self-Service user that the data in one timeline segment is a completely new data start. In other words, while the data within one segment is logically contiguous, the data is never contiguous across segments. For example, the following image shows a timeline with multiple segments.



Segmented Branch Timeline

As mentioned above, the branch timeline becomes segmented after you have performed a specific action or task, such as **Refresh**. Based on the action, two red bubbles will appear in the time segment. The top bubble indicates where the data used for this action came from – for example, the data template, a different branch, or a shared bookmark. The second red bubble appears on the timeline as the actual data stream in a point in time from the parent data. It appears because of actions such as **Refresh**, **Reset**, **Restore**, **Create Branch** and **Bookmark**. Clicking the second bubble will show you specific details of the action, such as the specifics of the action including its name, the time the action occurred, and the data sources used at a point in time. This is illustrated below.




Parent Data Sources and Child Data Sources

Working with Multiple Branches and Timelines

As you work in your data container, you can switch between branches to work on resolving a bug or to test a new application feature.


For example, consider what occurs on two different branches in a container:

Branch 1:




Branch 1 Timeline

Branch 2:




Branch 2 Timeline



A user may have actually worked with these branches in the following order over time:

Branch 1: Create a branch and use




Branch 2: Create another branch and use



Branch 1: Activate branch, Restore the data source and use




Branch 2: Activate branch and create bookmarks




Branch 2: Refresh the data source from a particular point in time



Branch 2: Reset a branch to the last action (e.g., refresh) on the timeline, and use



In the above illustrations, an individual branch's timeline shows all actions performed on the branch while the branch was active. The active branch timeline can be interrupted and deactivated when you choose to perform actions such as switching to another branch, **Create Branch**, **Activate**, or **Stop** a data container. Additionally, you will only be able to view actions on a single branch at a time. A better way to manage multiple branches is to go to the **Time** tab in the **Data Container View Panel**. The **Time** tab allows you to access the **container timeline**, which becomes useful as you toggle back and forth between branches to complete different tasks. The **container timeline** allows you to view all the continuous data points of time, with all actions taken on all branches in a single data container.




Container Timeline

 Merging keystores requires several manual steps. Refer to [Migrating a TDE-enabled vPDB](#) for more information.

Understanding How to Preserve Data in a Point in Time

The following illustration shows that on 8/27/14, at 9:33:09am, data was reset to the parent data branch (master) at 9:28:48am, capturing data points from 9:33:06am.




Preserve Data in a Point in Time

The black arrows above point to a tick, (representing a point in time) clicked on the branch timeline. This represents the time the **Reset** action was performed on the data container. The red arrows point to when time was captured in a data source using the **Reset** action on the branch timeline. When clicked, the reset bubble provides more details with a flyout, indicating where the data comes from and the time that the data represents. Additionally, the reset

bubble detail flip card provides additional information about each data source. Specifically, the blue arrows point to the time used for each data source at this point in the data container.

 This does not show the time that was used for each source that pulled the data.

The time represented on the branch timeline varies based on many factors. For example, after you select a specific point in time on the branch timeline, the Delphix Engine will map that point to the closest usable point in time for each data source. Based on the properties of the underlying data sources, these times may be different. Not all data sources track changes at the same granularity, as illustrated below.



Point In Time

While a branch timeline can follow a continuous-time flow, the data sources being selected for each time segment may not be continuous.

Understanding Bookmarks

Bookmarks are a way to mark and name a particular moment of data on a timeline that makes it easy to search and find events later. Once a point in time is selected, you can create a bookmark. You can select a specific point in time on the branch timeline, the Delphix Engine will then map that point to the closest usable point in time for each data source. Based on the properties of the underlying data sources, time may not be the same for every data source in the container. Not all data sources track changes at the same granularity. For more information, see [Working with Bookmarks in a Data Container](#).


 **Note**

A bookmark captures the latest point-in-time (not exceeding the bookmark timestamp) that is available while creating the bookmark. When you access the bookmark later, it does not get a closer time even if new logs are available.

Once created, you can easily locate a bookmark through one of the bookmark viewers in the interface.

Bookmarks Tab in the Data Container View Panel


The **Bookmarks** tab is the third tab in the **Data Container View Panel** within the data container workspace of the Delphix Self-Service interface. It allows you to find a bookmark that is within your data container and view the branch where the bookmark has been placed.



Bookmark Tab

Bookmarks Tile in the Data Container Report Panel

The **Bookmarks** tile in the **Data Container Report Panel** allows you to see all bookmarks within your container and all bookmarks that other users have made available to you. Here you can also edit details about bookmarks, create new branches, and restore the active branch to the bookmark's point of data time.



Bookmark Tile

Bookmark Sharing Permissions


When you first create them, bookmarks are private to your data container, but you can share a bookmark with other data users. Bookmarks that other users have shared with you are called "available" bookmarks. Your bookmarks will only be shared with data users in data containers created from the same data template. This is because all data containers created from the same data template have a compatible set of data sources.

Bookmark Appearance

	<p>A bookmark that is private</p>
	<p>A bookmark you have shared</p>
	<p>A bookmark that has been shared with you</p>

Data Container Storage and Retention for Branches and Timelines

Bookmarks mark a moment of data. Delphix Self-Service will never automatically delete the data marked by a bookmark. However, Delphix Self-Service will automatically delete a bookmark with an expiration date set after it has expired. For more information on setting or removing an expiration date, see [Data Container Activities](#). Delphix Self-Service may delete data from any time in the past on your branches, depending on the retention policies configured by your administrator. If you select a moment of data that has been deleted, the flyout will indicate that retention has removed data for this point in time.



Data Container and Retention

Data Container Activities

Getting Started

Data containers can be shared between multiple users. In this situation, users should coordinate with their co-owners when performing data operations that could disrupt other users' workflow such as stopping or refreshing the data container.

Activity One: How to Start and Stop a Data Container

Starting a data container does the following:

- Starts the data sources
 - This means that each data source listed in the **Source Details** section of the **Data Container** page will start using CPU and network resources on the target system it is running on
- Makes the data in the active branch available
 - Once the container has been started, the data represented by the active branch is available

Stopping a Data Container does the following:

- Shuts down the data sources
 - This means each data source listed in the **Source Details** section of the **Data Container** page will stop using CPU and network resources on the target system.

To start a data container, click **Start** on the **Self-Service Toolbar**.

To stop a data container, click **Stop** on the **Self-Service Toolbar**.

Working with a Branch, a Branch Timeline, and the Self-Service Toolbar

Activity Two: Using Reset from a Bookmark to Facilitate Destructive Testing

Reset is a data user workflow that is optimized to enable destructive testing. **Reset** automatically restores the data to the last operation conducted in the data container, which can include creating a bookmark, resetting, or restoring data. As an example, you can do a refresh and then get your data into a state required for testing. Once you are satisfied with the state of your data, you can create a bookmark, which will preserve the data at this point in time.

Afterward, you can then run destructive tests on the data. When you are done, you can click the **Reset** icon, which will automatically restore the state of the container to the last operation – in this case, the bookmark. This workflow ensures that each test has a clean copy of the data and is not impacted by the results of other tests. You only need to create a bookmark and click **Reset** on the **Self-Service Toolbar**.

Create a Bookmark

1. Select a **Data Point** on a branch's timeline.
2. On the **Self-Service Toolbar**, click the **Bookmark** icon.
3. In the **Bookmark Window**, enter a new **name**.

Create Bookmark

New bookmark

Created At default
Fri, Jul 28, 2017 15:29:52 PM

Description

Auto-Deletion
 Will be Deleted After
Jul 31, 2017

Tags **ADD A TAG +**
None

CANCEL **CREATE**

Name a Created Bookmark


4. Optionally, fill in a **description**.
5. Optionally, set an expiration date. The bookmark will be automatically deleted at the end of this day.
6. Optionally, add one or more **tags**.
You can use these to help filter a set of bookmarks.
7. Click **Create**.

After the bookmark has been created, you will see the **bookmark** icon appear on the timeline. When you click the **Reset** button, all data will be reset to that point in time.

Reset to Data from a Bookmark

Click the **Reset** icon.

This action reflects the moment of data marked by the closest operation bubble (**Refresh**, **Restore**, **Reset**, or **Bookmark**) into a new timeline segment on the active branch. It also copies the moment of data into the data sources.




Update Data with Reset

Undo Data

Click the **Undo Data** icon for an operation.


Undo the given operation. This is only valid for RESET, RESTORE, UNDO, and REFRESH operations.



Activity Three: Using Refresh to Get the Latest Data From a Data Template

1. Start a new timeline segment with the most recent point of data from the data container's data template.
2. Click the **Refresh** icon.

Refresh creates a new timeline segment on the active branch. This refreshes each source in the data container to the latest data in the corresponding source of the data template.



Update Data with Refresh

Activity Four: Using Restore to Return Data Back to a Point in Time

This starts a new timeline segment on the active branch with the selected point of data.

1. Select one of the following:
 - a. A **point of data** on a timeline.
 - b. A **bookmark** on a timeline.
 - c. A **bookmark** under the **Bookmarks** tile in the **Data Container Report Panel**.
2. Click the **Restore** icon.


If you restore data back to a point in time on the data template master timeline, you will be asked which data container to restore into. It will then:

- Reflect the selected point of data into a new timeline segment on the active branch
- Copy the moment of data into the data sources

If the timeline segment on a branch timeline was created by a **Restore** operation, then the segment starts with the moment of data from the branch that was selected when the **Restore** operation was done. This is illustrated below.

i Parent Branch

The parent branch for this segment can be the same branch of which this segment is a part. It is possible to restore the active branch from a point in time on the same branch.



Update Data with Restore

✓ Source Branch

The source branch for this segment can be the same branch of which this segment is a part. It is possible to restore the active branch from a point in time on the same branch.

Activity Five: Restoring to a Point on the Parent Template

Data templates serve as the parent for a set of data containers, and as a data user, you have the flexibility to restore your container to any point on the template.

1. Choose the container tab and select the template by clicking on the template name located at the top of the screen.
2. Select one of the following:
 - A point of data on the timeline
 - A bookmark on the timeline
 - A bookmark under the Bookmarks tile in the Data Container Report Panel
3. Click the restore icon
4. A dialog will pop up. Use it to select the container you'd like to restore.

forceOption

forceOption is an API/CLI-only feature. Generally, if a source database is corrupted or otherwise prevents taking VDB snapshots, the Refresh, Reset, and Restore actions cannot be completed. With **forceOption**, you can bypass taking a pre-operation screenshot and proceed with the desired action.

Because **forceOption** does not take a snapshot of the VDB before refreshing/resetting/restoring, you cannot undo the operation afterward.

From the CLI:

1. Go to the Delphix Self-Service container endpoint.
2. Select the container.
3. Attempt to execute the operation.
4. Set **forceOption** parameter to **true**.
5. Commit the change.

The operation will now perform without taking a snapshot first.

Activity Six: Create a New Branch and Switch Between Branches


Developers and QA teams can have multiple branches that can represent data from different points in time or different sources. You have many options for how you create a new branch. These include:

- A **point of data** time on a data timeline within the Delphix Self-Service data container, or
- A **bookmark** bubble on the timeline, or
- A **bookmark** in the **Bookmarks** tile in the **Data Container Report Panel**

1. Click the **Branch**  icon to create a new branch.
2. Enter a **name** for the new branch.
3. Click **OK**.
4. On the **Self-Service Toolbar**, click the **Activate** icon. 

If the inactive branch is not showing in the data container workspace:

1. Find the **branch** in the **Branch** tab.



Selection of Branches in Branch Tab

2. Click the **Activate**  icon.
3. After a moment, the branch will become active.

Active Branch

Within a single data container, only one branch is active at any given time. The data located at the red star of the active branch's timeline is the newest copy of the data from the data container's data sources. The active branch is distinguished by a red star, which appears at the far right of the timeline, alongside its name in the **Branch Name** area, and in the **Branch** tab.

Active branch	Inactive branch

Activity Seven: Rename and/or Delete a Branch

Rename the Default Branch

1. Select the **Default Branch** in the **Branch** tab.
2. Click the **Pencil** icon to the right of the name.
3. Enter the **new name**.
4. Click the **Checkmark** icon.

Delete a Created Branch

1. Select the **branch** in the **Branch** tab.
2. Click the **Delete** icon to the right of the name. 
3. Click **Delete** in the confirmation window that appears.

Activity Eight: Restoring a Data Container to a Consistent State with the Recovery Operation


Data Containers Consistency

Delphix Self-Service allows you to group multiple datasets in the same data container. This makes it easy for you to access entire applications such as PeopleSoft, including binaries and code. If a data container represents an application, then there are likely to be dependencies between the application's datasets. For example, the vFiles data source containing the code will depend on a specific version of the database's schema. Therefore, it is important that all dataset sources are drawn from the same point in time. If they are, the data container is in a "consistent" state; if they are out of sync, or "inconsistent," errors will occur. For example, if the vFiles data source containing the code has been updated more recently than the database's schema, the dependency cannot work.

Delphix Self-Service currently has no way to determine whether the application is consistent. However, it attempts to minimize the chance that dataset sources are out of sync whenever it performs a data operation such as refresh, restore, or reset. When performing a data operation, Delphix Self-Service attempts to snapshot all dataset sources from a point in time as close as possible to the desired time. If at least one of the data sources fails to go to the desired point, then Delphix Self-Service considers the data container to be in an inconsistent state. The application as a whole may still be working, but Delphix Self-Service assumes that the failed dataset's data is not the correct version. To return to a consistent state, you must perform a recovery operation on the data container.

Data Container Recovery

Prior to performing any data operation, Delphix Self-Service takes snapshots of all datasets. Recovery is the process of rolling back a data container to a snapshot, thereby restoring it to a consistent state. When a failure occurs, you will see the following screen:




You can either perform recovery or use a different data container. Whether the recovery will fail or succeed depends on exactly why the data operation failed in the first place. If the problem was intermittent, such as a temporary network problem causing SSH failure, then performing recovery should work. If the problem is persistent – for example, the target host is out of space – then intervention is required; recovery will not succeed until you address the underlying root cause of the failure.


Admins can see the underlying failure in the **Actions** sidebar or the **Job History** dashboard. The **Actions** sidebar is the preferred place to view the failure; it has a hierarchical display that makes diagnosing the failure more straightforward.

Activity Nine: Working with Container Locks

Container owners have the ability to lock/unlock containers. By default, a container is unlocked, which means that all the container's owners can perform operations on it:



You can lock a container by clicking the **Lock** button in order to prevent other users from operating on it.




The information on who locked the container is available in three places:

- In the history below the timeline
- In the tooltip on the locked container's name in the container selector
- In the tooltip on the locked container's name in the container label above the timeline

Errors

If an action such as refresh fails, a dialog box will give you further information. The **Error Details** field will

The **copy** button enables you to copy and paste error details, which you can then send to your Delphix administrator for further assistance.



Error dialog

Containers with Multiple Owners

Delphix Self-Service administrators can designate multiple users as owners of a single data container. These users all share access to the same data container which means actions taken by one user will impact all users on the same data container. For example, if User A activates Branch X, User B will also see Branch X as the active branch. This ability to for one user's actions to impact another user on the same containers creates new concerns for users sharing the same container. As a result, more processes should be put into place in order to coordinate usage between users. Each team is different, but strategies include:

- designating a person to perform certain data operations
- saving your work with a bookmark or creating / working on a personal branch
- being aware of who is using your data container / data before performing operations
- locking a container to prevent others from performing any operations on it

How many owners should a container ideally be shared between?

There is no technical limit built into the software, but it is best if a team of 5-10 users shares a single data container. In most cases, having fewer owners minimizes overhead and conflicting usage. One owner per container provides maximum productivity and minimal overhead, so this feature should only be used if your infrastructure or processes require that multiple users share a container. Additionally, Jet Stream Only users currently cannot see other users with whom they share the container.

How should users handle potentially disruptive operations?

If one user performs an operation on a data container, it will affect the other owners of that container. Additionally, each user has permission to perform the same operations on the data container; currently there are no fine-granularity permissions that limit the operations a user can perform. All operations are potentially disruptive, but the level of disruption varies by operation. If any of the following operations are performed at the same time, the second operation will fail due to a conflict when processing the job.

Conflicting operations

- Refresh
- Restore
- Reset
- Enable/Disable
- Create Branch
- Activate Branch
- Delete Branch
- Create Bookmark
- Delete Bookmark

If User A performs a destructive operation while User B is "using" the data container, the operation will destroy User B's current state. Currently, the interface does not provide insight into whether the data container is in use by another user.

Destructive operations


- Refresh
- Restore
- Reset
- Enable/Disable
- Create Branch
- Activate Branch

Deleting objects

All owners can delete any bookmarks or branches in the container, regardless of who created them.

Working with Bookmarks in a Data Container

Working with bookmarks is an easy way to share data with other users of any container created from the same template. By sharing with others, you can integrate testing, development, and QA needs. For example, in the past, if you found a bug you would wait until it was fixed. But with bookmarks, you do not have to stop your work while someone tries to fix the problem. Sharing a bookmark allows users to work with data as they see fit. Bookmarks mark a moment of data. Delphix Self-Service will never automatically delete the data marked by that bookmark. However, Delphix Self-Service will automatically delete a bookmark with an expiration date set at the end of that day.



Bookmarks Management in the Data Report View Panel

Activity Nine: Share a Bookmark with Other Delphix Self-Service Users

Share a Bookmark

1. Select a **bookmark** by clicking one of the following:
 - The bookmark's **bubble** on the **branch timeline**.
 - The **Bookmarks** tab in the data container workspace.
 - The **Bookmarks** tile in the **Data Container Report Panel**.

2. Click the **Share**  icon.

You cannot share a bookmark that you or another user have already shared.

Unshare a Bookmark

1. Select a **bookmark** by clicking one of the following:
 - The bookmark's **bubble** on the **branch timeline**.
 - The **Bookmarks** tab in the data container workspace.
 - The **Bookmarks** tile in the **Data Container Report Panel**.

2. Click the **Unshare**  icon.

You cannot unshare a bookmark that is already private or a bookmark which someone else has shared.

Delete a Bookmark

1. Select a **bookmark** by clicking one of the following:
 - The bookmark's **bubble** on the **branch timeline**.
 - The **Bookmarks** tile in the **Data Container Report Panel**.
2. Click the **Delete**  icon.

Activity Ten: Editing Bookmarks

Rename a Bookmark

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
A selection of bookmarks will appear based on whether you have chosen to view private, shared, and/or available bookmarks.
2. In the **detail bookmarks** window, click the **Edit** icon to the right of its name.
3. Enter the **new name** in the edit field.
4. Click the **checkmark** to the right of the field to accept and save the new name.

Edit the Description of a Bookmark

1. Select a **bookmark** by clicking the **Bookmarks** tile in the **Data Container Report Panel**.
2. Click the **Edit** icon to the right of its name.

Remove the expiration date of a Bookmark

1. Select a **bookmark** by clicking the **Bookmarks** tile in the **Data Container Report Panel**.
2. To the right of the bookmark's name, click the **Edit** icon.
3. Uncheck the **Will be deleted after** checkbox.
4. Click the **checkmark** to the right of the date selector.

Set or update the expiration date of a Bookmark

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
2. To the right of the bookmark's name, click the **Edit** icon.
3. Check the **Will be deleted after** checkbox.
4. Use the date selector to pick a new date.
5. Click the **checkmark** to the right of the date selector.

Activity Eleven: Filter and View Bookmarks

View Only Your Created Bookmarks

In the **Bookmarks** tile in the **Data Container Report Panel**, bookmarks that belong to you are shown. To see only your own bookmarks:

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
2. De-select **Available**.

View Bookmarks You Have Shared with Others

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
2. De-select **Private**.

3. De-select **Available**.

Only your shared bookmarks will be shown.

View Bookmarks That Others Have Shared with You

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
2. De-select **Private**.
3. De-select **Shared**.
4. Select **Available**.

These are the bookmarks that have been shared with you.

Adding Tags To Your Bookmark

1. In the **Data Container Report Panel**, click the **Bookmarks** tile.
2. Select the **bookmark** to which you want to add tags.
3. Click **Add a Tag**.
4. Enter the **tag name**.
5. Click the **Accept** icon.

Your tags will be shown at the bottom of the **Bookmarks** tile in the **Data Container Report Panel**.

You can only add tags to bookmarks that you have created.

Finding Bookmarks

In either the **Bookmarks** tab in the data container workspace or the **Bookmarks** tile in the **Data Container Report Panel**:

- Type into the **Filter** field.

This will only show bookmarks that have names or tags that match the text you have entered.

Understanding Delphix Self-Service Usage

Usage Management Dashboard Overview

Data templates are comprised of dSources, virtual databases (VDBs), and vFiles. These data sources are controlled by the standard policies configured in the **Management** application of the Delphix Engine. As with existing containers, space will be reclaimed by the retention policy over time. As retention cleans up historical data, users will no longer be able to use those points in time to restore or branch. In Delphix Self-Service, an admin can create a bookmark on the data template timeline, which will prevent retention from cleaning up the data that bookmark references.

Data containers are comprised of VDBs provisioned from the sources defined in the data template. Similar to VDBs in the **Management** application, data containers' VDBs will share blocks with the source from which they are provisioned. This prevents the referenced data on the source from being cleaned up by retention. Retention for these VDBs is controlled by the standard Delphix Engine retention policies. As on templates, bookmarks in data containers will prevent storage from being reclaimed by retention. In addition, Self-Service will ensure that the latest data on each branch is never removed.

The **Usage** pages of the data templates and data containers provide information that can help you understand how storage is being used, how to reclaim space, and how much space you are able to reclaim.


Usage Overview is a top-level page, along with the **Data Operations** and **Overview** pages. It contains the space usage breakdowns by data templates and users.

Container Usage Overview

The **Usage Details** page, shows the space used by data containers provisioned from the template and the bookmarks created on the template.



The stacked bar graph shows information about the top 10 space users. You can re-sort the graph based on the fields in the **Sort by** legend on the top right-hand corner of the screen as seen in the image above. For example, if you want to know which data containers are sharing the most data with others, you can un-select **Shared (others data)** and **Unique** by clicking them in the legend.



i When the legend items are not selected, their corresponding colored boxes turn gray and the data is removed from the chart. The data and name will reappear when you re-select by click on the preferred grayed-out category.


The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this data container. This assumes that also delete underlying data sources.
- **Shared (others data)** – The amount of space that cannot be freed on the parent data template (or sibling data containers) because it is also being referenced by this data container due to Restore or Create Branch operations. The snapshots on the template or sibling container are what use up space.

- **Shared (self data)** – The amount of space that cannot be freed on this data container because it is also being referenced by sibling data containers due to Restore or Create Branch operations, via shared bookmarks
- **Unvirtualized** – The amount of space that would be used by the data in this container without Delphix virtualization

Bookmarks Usage Overview

As shown in the image above, the **Container Usage** page provides the usage information about bookmarks created on a template. The primary categories of information include **Unique**, **Shared (others data)**, and **Shared (self data)**.



The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this bookmark

- **Shared** – The amount of space referenced by this bookmark that cannot be freed by deleting this bookmark because it is also referenced by neighboring bookmarks or branches that have been created or restored from this bookmark
- **Externally Referenced** – The amount of space referenced by this bookmark cannot be freed by deleting it because it is also being referenced outside of Self-Service – for example, by a retention policy.

Branches Usage Overview

As detailed in the image above, the **Container Usage Details** page shows the usage information about the branches and bookmarks created on a container. The primary categories of information include **Unique**, **Shared (others data)**, and **Shared (self data)**.



The field categories display the following information:

- **Unique** – The amount of space that will be freed if you delete this branch
- **Shared (others data)** – The amount of space that cannot be freed on the parent data template or sibling branches because it is also being referenced by this branch due to Restore or Create Branch operations. The snapshots on the template or sibling container are what use up the space.

- **Shared (self data)** – The amount of space that cannot be freed on this branch because it is also being referenced by sibling data containers due to Restore or Create Branch operations, via shared bookmarks.

Related Topics

- [Understanding Delphix Self-Service Usage Management](#)

Delphix Masking

Click the button below to access the Masking Documentation site.

 Unknown macro: 'landbox'

Developer's Guide



Unknown macro: 'navbox'

Command Line Interface Guide

This section covers the following topics:

- [Command Line Interface Overview](#)
- [Delphix Objects](#)
- [Command Reference](#)
- [CLI Cookbook: Common Workflows, Tasks, and Examples](#)

Command Line Interface Overview

This topic provides an overview of the Delphix Engine command-line interface and links to additional topics.

The Delphix Engine provides a native command-line interface (CLI) accessible over SSH. This CLI provides an interactive layer on top of the public web service APIs, and is intended for users that wish to automate interactions with the Delphix Engine, or simply prefer a text-based interface. All of the functionality available in the CLI is also available through the public stable web service APIs should more full-featured automation be required. For more information on automation using the web service APIs, see the [Web Service API Guide](#).

The CLI has an internal help system and supports tab-completion to help guide users. Running the `help` command will display a list of valid commands and properties, if applicable. Specifying the command or property as an argument to `help` will display more specific information about that command or property. This guide serves as an overview of CLI operation and examples of some basic tasks, and is not a reference for all CLI commands or properties. As the CLI content is identical to the public web services, complete information about particular commands, properties, or other operations can be found in the API documentation delivered with each server instance, found at:

```
http://<server>/api
```

The API documentation is guaranteed to be consistent with the set of APIs exported by that particular server. All of the APIs used by the GUI will be supported by the CLI. While all the database and environment APIs are available, most of the system-oriented APIs (such as those required to do initial setup) will be made available in a later release.

- [Connecting to the CLI](#)
- [CLI Contexts](#)
- [Managing Objects](#)
- [Managing Properties](#)
- [Array Properties](#)
- [Untyped Object Properties](#)
- [CLI Automation](#)

Connecting to the CLI

This topic describes how to connect to the Delphix Engine command line interface.

 If the same username exists in both the SYSTEM and DOMAIN namespaces (for example a sysadmin and a domain user), when logging in via ssh it is necessary for the user to explicitly provide the namespace SYSTEM or DOMAIN. For example, a user whose name exists in both namespaces must ssh in as either "someone@SYSTEM" or "someone@DOMAIN". If the user tries to log in as just "someone", the engine will give an error:

Ambiguous username. Add a @DOMAIN or @SYSTEM suffix to the username.

The CLI is available over SSH or the terminal console on any Delphix Engine version 3.0 or later. To connect, use any SSH client appropriate for your workstation environment and connect to the Delphix Engine by IP or hostname on the standard SSH port (22). Enter a username for either a domain or system user followed by the namespace appropriate to that user (either DOMAIN or SYSTEM). For example:

- `ssh admin@DOMAIN@delphix-server.example.com`
- `ssh sysadmin@SYSTEM@delphix-server.example.com`

 The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

At the prompt, enter your user password. Once connected, you will be placed at the CLI prompt:

```
delphix>
```

While both `admin` and `sysadmin` produce the same prompt once logged in, be aware that the two users have different menus and different functional areas.

Sysadmin Menu

```
delphix> ls
Children
network
service
storage
system
user
Operations
version
Operations
version
```

```
delphix>
```

Delphix Admin Menu

```
delphix> ls
Children
alert
audit
authorization
connectivity
database
environment
fault
group
host
job
namespace
network
policy
replication
repository
service
session
snapshot
source
sourceconfig
system
timeflow
user
Operations
version
delphix>
```

Individual commands passed as arguments to the SSH client will be interpreted as if they had been read from the terminal. More complex scripts can be passed as input to the SSH command. When running SSH in non-interactive mode via these mechanisms, the command line prompt will be suppressed, as will terminal font decorations such as underline and bold.

The CLI is also available from the serial terminal console should the network be unavailable. Consult your VM platform documentation for information on how to connect to the terminal console. Once connected, log in using your Delphix user credentials just as you would over SSH.

If the management service is unavailable due to a software bug or other problem, the CLI can still be accessed as a system user provided that user is locally authenticated (not via LDAP) and has logged in at least once before. While in this state, only the `system` commands are available, including `restart`, which will attempt to restart the management service without rebooting the entire server. If this problem persists, please contact Delphix support.

The topic [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#) shows an example of how to connect to the CLI using SSH key exchange instead of the standard password-based authentication.

CLI Contexts

This topic explains the concept of contexts within the Delphix Engine command-line interface.

The CLI is built on the concept of modal “contexts” that represent an administrative point for interacting with the web service APIs. These contexts can be divided into the following types:

Context	Description
Static Children	These contexts exist for the purpose of navigating between points in the hierarchy, but have no properties of their own and do not correspond to any server-side object. The root context is an example of this, as are most of the top-level contexts such as <code>database</code> or <code>group</code> .
Object	These contexts represent an object on the server, either a specific object (such as databases) or system-wide state (such as SMTP configuration). These contexts have properties that can be retrieved via the <code>get</code> command.
Operation	These contexts represent a request to the server. Commands may or may not require input and may or may not change state on the server, but in all cases require an explicit commit operation to execute the command. When in a command context, the prompt includes a trailing asterisk (*) to indicate that <code>commit</code> or <code>discard</code> is required before exiting the context.

User can move between contents by typing the name of the context. To move to a previous context, the `up` or `back` commands can be used. In addition, the CLI supports UNIX-like aliases for `cd` and `ls`, allowing navigation similar to a UNIX filesystem. For more information on these commands, see the [Command Reference](#) section.

Managing Objects

This topic describes the use of objects in the Delphix Engine command-line interface and provides a list of the object management operations.

The Delphix Engine represents state through objects. These objects are typically managed through the following operations, covered in more detail in the [Command Reference](#) topics

The topic [CLI Cookbook: Changing the Default Group Name](#) illustrates the use of object management commands such as `list` and `get`.

Operation	Description
<code>list</code>	For a given object type (represented by a static context such as <code>database</code>), list the objects on the system, optionally constrained by some set of attributes. Some objects are global to the system and do not support this operation.
<code>select</code>	Select a particular object by name to get properties or perform an operation on the object. See the “Delphix Objects” section for more information on object naming.
<code>get</code>	Display all or some of the properties of an object after selecting it.
<code>update</code>	Enter a command context to change one or more properties of an object after selecting. Not all objects support this operation, and only properties that can be edited are shown when in the update command context.
<code>create</code>	Create a new instance of the object type from the root static context. Not all objects can be created in this simplified fashion. Databases, for example, are created through the <code>link</code> and <code>provision</code> commands.
<code>delete</code>	Deletes an object that has been selected. Not all objects can be deleted.

In contexts where there are multiple objects of a given type, the `list` command can be used to display available objects, and the `select` command can select an object for subsequent operation.

When listing objects, each context has its own set of default columns to display. The `display` option can be used to control what columns are displayed to the user. This is a comma-separated list of property names as they would be retrieved by the `get` command. It is possible to specify properties that do not exist in order to accommodate lists of objects of varying types and untyped objects.

The topic [CLI Cookbook: Listing Data Source Sizes](#) provide an example of using the `list` command.

Managing Properties

This topic describes the use of properties in relation to objects in the Delphix Engine command-line interface.

Object properties are represented as a hierarchy of typed name/value pairs. The `get` command by itself will display the complete hierarchy for a particular object. This hierarchy is displayed with each nested object indented by an additional level. The set of available properties depends on the command context and may change if the type of an object is changed.

Property State

Properties are typically set to a specific value, but they can also be `unset`. Unset properties indicate there is no known value, either because it hasn't been provided yet, or it has been explicitly removed. Properties in this state are displayed via the following means:

- `(unset)` – The property is not currently set. It may never have been given a value or it may have been explicitly unset through the `unset` command.
- `(required)` – This has the same underlying semantics as `(unset)`, but indicates that the property **must** be set before the current command can be committed. Failure to do so will result in a validation error at the time the commit operation is attempted. Required properties are displayed in bold.

In addition, all objects have a default state when in command context. A property that has been modified is noted with an asterisk (*), and can be reverted to its default state through the `revert` command.

When updating properties, only those properties are sent to the server. The exception is arrays and untyped objects, covered in [Array Properties](#) and [Untyped Object Properties](#). These objects are always sent in their entirety, so changing any one element will send the entire object.

Basic Properties

Most properties are displayed and input as a string, though the underlying type may be more specific. The following are some of the basic types:

- **String** – An arbitrary string. This may be subject to additional validation (such as an IP address) that is enforced at the time the property is set.
- **Number** – An integer number.
- **Boolean** – Either “true” or “false”.
- **Enumeration** – A string that must be chosen from a known set of options.

Nested Properties

Some properties are in fact other objects and are represented as a nested set of properties. These properties can be manipulated in one of two ways: by specifying a dot-delimited name or changing the context via the `edit` command.

A dot (.) in a property name indicates that the portion to the left of the dot is the parent object name, and the portion to the right is a child of that object. For example, `sourcingPolicy.logsSyncDisabled` denotes the `logsSyncDisabled` property within the `sourcingPolicy` property. These dots can be arbitrarily nested.

An alternative syntax of using brackets to enclose property names (`sourcingPolicy[logsSyncDisabled]`) is also supported for familiarity with other programming languages.

The `edit` command, in contrast, will change the current context such that all properties are relative to the specified object. This can be useful when changing many nested properties at once, or when the complete set of properties can be confusing to manage all at once.

The topic [CLI Cookbook: Disabling LogSync for a dSource](#) provides an example of manipulating nested properties.

Array Properties

This topic describes the use of array properties in the Delphix Engine command-line interface.

Some Delphix objects represent properties as arrays. Arrays are effectively objects whose namespace is a contiguous set of integers. While they behave like objects and their properties can be referenced via the same object property notation, they differ in several key areas.

Arrays can be divided into two types: arrays of primitive types (strings, integers, etc.) and arrays of objects. Arrays of objects can be managed like other objects via nested property names and the `edit` command, but differ in the following respects:

- When an array element is `unset`, it removes the element from the array and shifts all other elements down to preserve the contiguous index space.
- New array elements can only be appended to the end of the array by specifying an index that is one more than the maximum index of the array.
- When displaying a property that is an array, if the length is greater than 3, then it is displayed only as “[...]”. The complete contents of the array can be displayed by getting or editing that particular property.

Arrays of primitive types can be managed as arrays of objects, but also support an inline notation using comma-separated notation. This allows single-element arrays to be set as standard property, and for arrays of strings to be set on a single line instead of having to edit each element.

Regardless of the element type, arrays are sent as complete objects when updated. When any array element is changed and subsequently committed, the complete array is sent to the server. When a single array element is reverted, the entire contents of the array are reverted.

The topic [CLI Cookbook: Setting Multiple Addresses for a Target Host](#) provides an example of working with a property that is an array of strings.

Untyped Object Properties

This topic describes the use of the `type` field in the Delphix Engine command-line interface object model, and the use of untyped objects.,

Most Delphix objects are typed, meaning they have a `type` field that controls what properties are available and their types. Object types and their associated hierarchy are described in more detail in the topic [Object Type Hierarchy](#) topic. In contrast, some properties are “untyped” objects, which means that there are no constraints on the property namespace, and all properties are plain strings. These objects are used for database configuration templates and other scenarios where the property namespace is unbounded or under the control of the user.

Untyped objects are always sent in their entirety when making updates. This means that when any value is changed and then committed, all values are sent. In addition, when reverting a single value within an untyped object, the entire parent object is reverted to its default state.

CLI Automation

This topic describes using automation with both the Delphix Engine command-line interface (CLI) and the web service API.

All functionality is available in both because the CLI is built upon the web services API. The CLI enables you to create scripts for simple automation, and it is a useful aid in the development of more complex code that uses the web service API.

Using the CLI for Simple Scripts

For simple automation, you can build routines that make CLI calls through SSH.

This snippet lists all environment names. It leverages the SSH key exchange explained in [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#) so that no password is required for the user named "automation".

```
DELPHIX_ENGINE=172.16.180.33
SSH_CMD="ssh automation@${DELPHIX_ENGINE}"

${SSH_CMD} "cd host; list display=name"
```

Backward Compatibility

Both the CLI and web services API are versioned to support backward compatibility. Future Delphix versions are guaranteed to support clients that explicitly set a version provided the major version identifier is compatible. For more information, see the [Web Service API Guide](#). The CLI will always connect with the latest version, but the `version` command can be used to both display the current version and explicitly bind to a supported version.

Users building a stable set of scripts can run `version` to get the current version. Scripts can then run the `version <id>` command to guarantee that their scripts will be supported on future versions. For more information on the different API versions and how they map to Delphix versions, see the [API Version Information](#) section.

Parsing CLI Output

The default text output of the CLI is unstable. Any attempt to parse the output is certain to run into difficulties in repeatable results for unknown input, as well as instability as the text output is changed in subsequent releases. Column headings, column order, and the number of columns will change in subsequent releases.

You can specify a version in your scripts to counteract this, but you will not be able to take advantage of new features and fixes.

CLI as a Development Tool for Complex Automation

While the CLI is useful for simple automation tasks, it can be slow and overly complicated due to the many round trips needed to control the automation logic. For example, to disable all the environments for an engine, you could write a script which lists the environments and modifies each one:

```
DELPHIX_ENGINE=172.16.180.33
SSH_CMD="ssh automation@${DELPHIX_ENGINE}"
```

```
env_array=(` ${SSH_CMD} "version 1.5.0; cd environment; list display=name" | grep -v
NAME` )
for i in "${env_array[@]}"
do
  ${SSH_CMD} "version 1.5.0; cd environment; select $i; disable; commit"
done
```

This script works, but it will be slow on systems with many environments since each SSH command will start a new session.

The web service APIs are superior when performing many operations as a single logical unit. The web service APIs also provide substantially more data with a single call than what is shown in the CLI output, which can greatly simplify your code and avoid multiple round trips.

However, the input and output of web service API calls are JSON data, and it can be difficult to quickly determine what the input and output will look like.

For this reason, the CLI provides two options that can greatly assist you in the development of complex automation: **JSON Output** and **Tracing**.

`(setopt format=json)` changes the CLI to the output of all results to parseable JSON (javascript object notation). This is the fastest and easiest way to quickly see what the JSON output will look like when executed via the Web Service APIs. The JSON format has wide support in a variety of programming languages; see <http://www.json.org> for more information.

`(setopt trace=true)` will display the underlying HTTP calls being made with each operation and their JSON payload. This allows you to determine the GET and POST calls, and their JSON payloads, which perform the actions that you need to power your automation.

`(setopt format=text)` changes the CLI back into its regular output mode. `(setopt trace=false)` turns off the trace display.

 Using both options will show the JSON output twice

The fastest way to develop complex automation is to experiment with the CLI and copy the underlying API calls to a custom system for better control over behavior.

```
delphix421> setopt trace=true
delphix421> cd user
delphix421 user> create
delphix421 user create *> ls
Properties
  type: User
  name: (required)
  authenticationType: (unset)
  credential: (unset)
..... (Output Truncated) .....
  userType: DOMAIN
  workPhoneNumber: (unset)
delphix421 user create *> set name=Jose
delphix421 user create *> set authenticationType=NATIVE
delphix421 user create *> set credential.password=Password1
delphix421 user create *> commit;
```

```
=== POST /resources/json/delphix/user ===  
{  
  "type": "User",  
  "name": "Jose",  
  "authenticationType": "NATIVE",  
  "credential": {  
    "type": "PasswordCredential",  
    "password": "Password1"  
  }  
}  
=== RESPONSE ===  
{  
  "type": "OKResult",  
  "status": "OK",  
  "result": "USER-35",  
  "job": null,  
  "action": "ACTION-107"  
}
```

Using the output above, you can see that to create a user you must use the URL "http://myengine/resources/json/delphix/user". You will use a POST command and pass a JSON payload which looks like the above. You will get a JSON response like the above, and can validate that the status is "OK".

Delphix Objects

These topics describe the object model for the Delphix Engine command-line interface.

The Delphix object model is a flexible system for describing arbitrary hierarchies and relationships of objects. In order to enable the current and future functionality of the system, the relationship between objects is not always immediately obvious. The CLI is merely a veneer atop the web services layer to ensure that the full complement of functionality expressed by the API is always available, but this requires users to have some understanding of how objects are represented in the system.

This section covers the following topics:

- [Object Type Hierarchy](#)
- [Object Names and References](#)
- [Databases and Environments](#)
- [Asynchronous Jobs](#)

Object Type Hierarchy

This topic describes the object type hierarchy for the Delphix Engine command-line interface.

All Delphix objects have an associated type. This type determines what properties are available for a particular object, the format of those properties, and controls how the system interprets objects and commands. The type hierarchy uses polymorphic inheritance to allow for common properties and behavior to be defined at a single point while permitting dramatically different types of objects to co-exist without requiring a completely separate API for each. For example, the `SourceConfig` object is the base type for all external database configurations, but it has children that include `OracleSIConfig` and `OracleRACConfig` types that refer to a single instance and RAC databases, respectively.

When specifying input types, the system will attempt to determine types appropriate for the current operation, but there are times when the type must be explicitly set, either because the operation supports multiple possible inputs, or the object can embed an abstract type. In these cases, it may be necessary to explicitly set the `type` property. Setting the type may change the set of visible properties and the resulting validation that is performed, but it will not affect any properties that are already set.

Object Names and References

This topic describes the use of object names and references in the Delphix Engine command-line interface.

Most Delphix objects are persistent objects in that they have a well-known identity on the server and associated persistent state. The exceptions are objects used only as input to other operations, or global objects that have a persistent state but don't require any explicit identity since they always exist.

Persistent objects have both a name and a reference. The reference is the canonical identifier for the object and remains valid even if the object is renamed on the server. It is an opaque token that should never be interpreted by the client; the format may change in future releases though backward compatibility with current references will be maintained. All web service APIs operate using references. References can be used in the CLI when selecting objects, but given that they are a programmatically generated internal concept, they are difficult for most users to use.

The object name, on the other hand, is a much more convenient way to refer to objects but suffers from the fact that it is not guaranteed to be globally unique. When displaying or setting references, the CLI will convert to or from the 'canonical name' based on the type of the reference and the current set of objects on the system. The canonical name has the form:

```
<Type>:/<Parent>/<Object>@<Namespace>
```

The type, parent, and namespace are only included if the local object name conflicts with other objects on the system that would otherwise be valid for the given type specification. Not all objects have names relative to their parent; groups, environments, users, and many other objects are globally unique on the system. This "best fit" method is used both when displaying references as well as when setting properties that are references. If the given name potentially matches multiple objects when attempting to set a reference property, then an error is displayed that includes a list of possible names to clarify which object is being referred to. The conversion from reference to name on display only happens with text output format. When the output format is JSON, the raw content is displayed (including the local name) and it is up to the consumer to format names appropriately based on their semantics. The conversion from name to reference when setting properties always occurs. Consumers can use references, optionally prefixed with a backtick (`) character to signify they are references in the unlikely event that someone has created an object with the same name as a valid reference.

- ✓ Providing unique names for objects without the use of forward slashes (/) and at signs (@) will provide the simplest CLI experience when referencing objects.

Here are some scenarios for databases and groups and their resulting behavior:

No conflicting database name

The local name will be used when displaying references to the object, and can be used when setting references:

```
set container=example
```

Databases with the same name in different groups

The parent group name must be used when displaying references to the object and when setting references to the object:

```
set container=group1/example
```

Databases with the same name in different namespaces

The namespace name must be used when displaying references to the object and when setting references to the object:

```
set container=example@namespace
```

Objects of different types but with the same name

This conflict is exceptionally rare, as the reference context typically constrains the set of possible objects to be a single type, but there are cases (such as alerts, or policy targets) that can be applied to any object. In these cases, the type name must be included to uniquely identify the object:

```
set target=Container:/group1/example
```

In the event that one of the named components contains a slash ('/') or an at-sign ('@'), single quotes must be used to disambiguate the name from its parent or namespace.

Databases and Environments

This topic describes the relationship between database container objects and environments in the Delphix Engine object model.

The core Delphix objects revolve around the notion of environments and databases, known at the API layer as **containers**. Understanding how these objects relate to each other is crucial to operating effectively within the CLI. This section provides an overview of these objects; for more information about a particular representation such as Oracle RAC, see the [Web Service API Guide](#).

i In variations of a use case where a database propagates down to lower environment databases (for keeping consistent data), records created in the lower tables would likely be overwritten by the propagating database once they are synced. Delphix is keeping track of changes at the data block level on disk, which would render this result.


As a solution, you could create a clone of the propagating database and then programmatically compare/merge the changes. You can also use a product like Oracle Goldengate, HVR, IBM Datastage, etc, to capture (extract) and propagate (replicate to lower) changes.

Environment Components

- An **environment** is the root representation of the external state that manages database instances. An environment could be a single host (`UnixHostEnvironment`) or an Oracle cluster (`OracleClusterEnvironment`). Environments exist to contain repositories, and each environment may have any number of repositories associated with it.
- A **repository** is an entity that contains database instances. Repositories are typically installation directories (`OracleInstall`) within an environment. Within each repository is any number of `SourceConfig` objects, which represent known database instances.
- The **source config** exists independent of Delphix, and could represent a possible dSource (in which case there is no associated database object), or could be managed entirely by Delphix (for VDBs). The source config contains intrinsic properties of the database instance, while the source (described below) contains information specific to Delphix and only exists when the source config is linked to a dSource or VDB.

Most environment objects are created through the act of **discovery**. By specifying a host, Delphix will attempt to automatically discover all environments, repositories, and source configs. These objects can also be added manually after the fact in cases where discovery fails.

The environment hierarchy can be represented as depicted below. The generic type is listed in the top portion of each box, with an example of the Oracle single instance objects in the lower portion of each box. Each of these objects can contain multiple child objects with it.




Database Components

The core of all databases within Delphix is the container that holds all the physical data associated with the database, whether a dSource or VDB. Within each container is a **Timeflow**, which represents a single timeline of change within the database history. Currently, a container can only have one Timeflow, though this limitation may be relaxed in a future release.

Within a Timeflow are two important objects: `TimeflowSnapshot` objects and `TimeflowRange` objects. Timeflow **range** represents the provisionable ranges within the history of the Timeflow, while Timeflow **snapshot** represents a point at which a snapshot was taken and therefore is more likely to provision in a short amount of time.

The environment hierarchy can be represented as depicted below. Each container may be associated with a **source**, which is the Delphix representation of an external database when it is associated with a container, and contains information specific to managing that source. Not all source configs within an environment have a source associated with them (as is the case with linkable databases), but all sources must have a source config. Containers may have no sources associated with them if they are unlined; sources can be manually attached at a later point. Currently, each container can have at most once source associated with it, though this may change in a future release.



Asynchronous Jobs

This topic describes conditions under which command-line interface operations may spawn jobs that run in the background and using the `wait` option to wait for job completion.

Not all operations can be performed in the context of a single web service API call. For cases where there is a long-running operation that cannot be executed quickly and transactionally, a job may be dispatched to do the remaining work in the background. For more information on jobs and their semantics, see the topic [Viewing Action Status](#). Within the CLI, any command can potentially result in an asynchronous operation. The default behavior is to wait for any such job to complete and display its progress in the CLI.

In the event that you do not want to wait for the operation to complete, the global wait option can be set (`setopt wait=false`). If disabled, the CLI will display the reference to any job that was dispatched, but not wait for it to complete.

Command Reference

These topics describe the core built-in commands within the CLI. It is not an exhaustive list of all commands in all contexts. For an object or type-specific commands, consult the API documentation.

This section covers the following topics:

- [CLI Help and Display Commands](#)
- [CLI Context Commands](#)
- [CLI Object Commands](#)
- [CLI Miscellaneous Commands](#)
- [CLI Property Commands](#)

CLI Help and Display Commands

This topic describes help and display commands for the Delphix Engine command-line interface.

Command	Description
<code>children</code>	Display all statically defined children valid for the current context. These children can be targets of the <code>cd</code> command.
<code>commands</code>	Display all build in commands valid for this context.
<code>help</code>	Display all commands and properties valid for the current context. Specifying a command or property will provide more information about that command or object. When nested properties are present, only top-level properties are displayed by default, though specifying a particular property will display the entire hierarchy.
<code>ls</code>	Display children, commands, objects, and operations valid in the current context. Only those sections that are relevant in the current context are displayed.
<code>operations</code>	Display available context-specific operations. These operations require an explicit <code>commit</code> command to execute the operation, or <code>discard</code> to abort it.

CLI Context Commands

This topic describes context commands for the Delphix Engine command-line interface.

Command	Description
<code>back</code>	Return to the previous visited valid context. This history only tracks contexts that were actually visited, so running <code>database "example"</code> followed by <code>back</code> will return you to the root context, not the database (because the two were executed as part of one action and never actually visited). If a previous context was deleted or is no longer valid, this command will skip over it.
<code>cd</code>	Switch to the given child. This is identical to typing the name of the child itself, but also support UNIX-style directory structures, such as <code>/</code> and <code>..</code> . This allows for contexts to be chained such as <code>cd ../database/template</code> .
<code>history</code>	Display the history of input to the shell. The shell supports the ability to move back and forth in the history using the up and down arrows.
<code>up</code>	This is an alias for <code>cd ..</code> for the benefit of those less familiar with UNIX filesystem navigation. Unlike <code>back</code> , which only returns to the previous context only if it was visited, and may return to a child context, this command will always return to the immediate parent context.

CLI Object Commands

This topic describes object commands for the Delphix Engine Command Line interface.

Command	Description
<code>list</code>	List all objects of a particular type when in the appropriate root context. Different contexts may support different options to the list command to constrain the output; run <code>help list</code> to see possibilities.
<code>select</code>	Select an object by name within a list.

CLI Miscellaneous Commands

This topic describes miscellaneous commands for the Delphix Engine command-line interface.

Command	Description
<code>echo</code>	Print the input arguments.
<code>exit</code>	Exit from the current CLI session. This is equivalent to sending the EOF control character (typically Ctrl-D) or closing your client SSH application.
<code>getopt</code>	Get the current value of a global configuration option. The list of global options can be retrieved by running <code>help getopt</code> , but include options for controlling JSON output (<code>format</code>), tracing HTTP calls (<code>trace</code>), and enabling synchronous job semantics (<code>wait</code>).
<code>setopt</code>	Set the value of a global configuration option.
<code>version</code>	Display the current API version or bind to a particular version. See the CLI Automation section for more information.

CLI Property Commands

This topic describes property commands for the DelphixEngine command-line interface.

Command	Description
<code>commit</code>	When in operation context, commit the changes and execute the operation.
<code>discard</code>	When in operation context, discard any changes and abort the operation.
<code>edit</code>	Change the current context to be relative to a particular object property when in operation context.
<code>get</code>	Get all properties (with no arguments) or a particular property of the current object.
<code>revert</code>	Revert a particular property to its default value, either the value of the underlying object during an update or the default command input value.
<code>set</code>	Set the value of one or more properties. These properties can be specified as <code>name=value</code> , or as simply the property name. When only the property name is specified the CLI will prompt for the value to use, optionally obscuring the input if the property is a password.
<code>unset</code>	Clear the current value of a property. This is not the same as reverting the property, though this can have semantically identical behavior in the case that the default value is unset.

CLI Cookbook: Common Workflows, Tasks, and Examples

This section covers the following topics:

- [CLI Cookbook: Authentication and Users](#)
- [CLI Cookbook: System Administration](#)
- [CLI Cookbook: Hosts and Environments](#)
- [CLI Cookbook: Source Databases and dSources](#)
- [CLI Cookbook: VDBs](#)
- [CLI Cookbook: Replication](#)
- [CLI Cookbook: Delphix Self-Service Actions](#)
- [CLI Cookbook: Hooks and Hook Templates](#)
- [CLI Cookbook: Network Performance](#)
- [Kerberos CLIs](#)

CLI Cookbook: Authentication and Users

These topics describe command-line interface procedures for authentication and managing users.

- [CLI Cookbook: Configuring SAP ASE Manual Discovery](#)
- [CLI Cookbook: Changing HTTP and HTTPS Web Connections](#)
- [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#)
- [CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users](#)
- [CLI Cookbook: Configuring SSH Host Verification for UNIX Environments](#)
- [CLI Cookbook: Password Policies](#)
- [CLI Cookbook: Changing the Default admin username](#)
- [CLI Cookbook: Disabling an admin User](#)

CLI Cookbook: Configuring SAP ASE Manual Discovery

Overview

This topic describes how to use CLI commands to manually add ASE repositories to an SAP ASE environment. Discovery is the process by which the Delphix Engine identifies data sources and data dependencies on a remote environment. SAP ASE repository discovery is done automatically when an environment is added to the Delphix Engine or when an already added environment is refreshed. In some cases, automatic discovery does not discover all of the repositories in an SAP ASE environment. These repositories may be added using manual discovery.

⚠ Unlike automatically discovered instances, manually discovered instances are not automatically deleted if the environment is refreshed when the instance is not running. The physical attributes such as ASE listener port, installation directory, and instance owner are not updated during an environment refresh either. If you change a physical attribute, you must manually update the repository.

To manually discover an SAP ASE repository you will need to:

- Add an SAP ASE environment
- Use CLI to manually discover a repository

Creating an SAP ASE Environment

Please refer to [Adding an SAP ASE Environment](#) for detailed steps.

Manually Discover a Repository

This example uses **sc-dev3.dc1** as the example environment.

1. Log into CLI and cd to the **repository** menu:

```
$ ssh admin@sc-dev3.dc1
Password:
sc-dev3.dc1> cd repository
sc-dev3.dc1 repository>
```

2. Add (manually discover) an SAP ASE repository instance:

⚠ The values used in the following code block are specific to the example instance we are adding.

```
sc-dev3.dc1 repository> create
sc-dev3.dc1 repository create *> ls
Properties
  type: ASEInstance
  credentials: (unset)
  dbUser: (unset)
  dumpHistoryFile: (unset)
```

```

environment: (required)
installationPath: (required)
instanceName: (required)
instanceOwner: (required)
isqlPath: (unset)
linkingEnabled: true
ports: (required)
provisioningEnabled: true
servicePrincipalName: (unset)
staging: false
version: (unset)
sc-dev3.dc1 repository create *> set credentials.type=PasswordCredential
sc-dev3.dc1 repository create *> set credentials.password=sybase
sc-dev3.dc1 repository create *> set dbUser=sa
sc-dev3.dc1 repository create *> set environment=rh610-ebf-ase-tgt-
dev3.dc3.delphix.com
sc-dev3.dc1 repository create *> set installationPath=/opt/sybase/15-7/sp139/install
sc-dev3.dc1 repository create *> set instanceName=ASE157_TGT
sc-dev3.dc1 repository create *> set instanceOwner=sybase
sc-dev3.dc1 repository create *> set ports=5400
sc-dev3.dc1 repository create *> ls
Properties
  type: ASEInstance
  credentials:
    type: PasswordCredential (*)
    password: ***** (*)
  dbUser: sa (*)
  dumpHistoryFile: (unset)
  environment: rh610-ebf-ase-tgt-dev3.dc3.delphix.com (*)
  installationPath: /opt/sybase/15-7/sp139/install (*)
  instanceName: ASE157_TGT (*)
  instanceOwner: sybase (*)
  isqlPath: (unset)
  linkingEnabled: true
  ports: 5400 (*)
  provisioningEnabled: true
  servicePrincipalName: (unset)
  staging: false
  version: (unset)
sc-dev3.dc1 repository create *> commit
`ASE_INSTANCE-3
sc-dev3.dc1 repository>

```

Updating a repository

Adding onto the above, the following example illustrates updating an SAP ASE instance's version after upgrading SAP ASE:

 Take caution when setting the version string. Make sure it matches the output as displayed by the "select @@version" query all the way out to the patch level (PL). For example, "15.7 SP138" or "16.0 SP02 PL01".

```
sc-dev3.dc1> repository
sc-dev3.dc1 repository> select ASE157_TGT
sc-dev3.dc1 repository 'ASE157_TGT'> update
sc-dev3.dc1 repository 'ASE157_TGT' update *> set version="15.7 SP139"
sc-dev3.dc1 repository 'ASE157_TGT' update *> ls
Properties
  type: ASEInstance
  credentials:
    type: PasswordCredential
    password: *****
  dbUser: sa
  dumpHistoryFile: (unset)
  installationPath: /opt/sybase/15-7/sp139/install
  instanceOwner: sybase
  isqlPath: /opt/sybase/15-7/sp139/install/OCS-15_0/bin/isql_r64
  linkingEnabled: true
  ports: 5400
  provisioningEnabled: true
  servicePrincipalName: (unset)
  staging: false
  version: 15.7 SP139
sc-dev3.dc1 repository 'ASE157_TGT' update *> commit
sc-dev3.dc1 repository 'ASE157_TGT'>
```

CLI Cookbook: Changing HTTP and HTTPS Web Connections

By default, the Delphix Engine allows both HTTP and HTTPS web connections. The following steps provide instructions on how to change their configuration:

1. Via CLI, login to the Delphix Engine as a system administrator (sysadmin).
2. Go to "service > httpConnector".
3. Initiate an update.
4. Set "httpMode" to the desired value among:
 - a. "BOTH": accepts HTTP and HTTPS connections (this is the default)
 - b. "HTTPS_ONLY": accepts only HTTPS connections
 - c. "HTTP_ONLY": accepts only HTTP connections
 - d. "HTTP_REDIRECT": accepts HTTPS connections and redirects HTTP connections to HTTPS.
 - e. "HTTP_REDIRECT_WITH_HSTS": redirect all requests made over HTTP to HTTPS and add Strict-Transport-Security Header to all responses.
5. Commit your Change. The Delphix web application will restart.

The following is an example of how to set the Delphix Engine to accept HTTPS connections and redirect HTTP connections to HTTPS.

```
cd /service/httpConnector
update
set httpMode=HTTP_REDIRECT
commit
```

Replacing the HTTPS (HTTP Secure) Certificate

This topic explains how to replace the HTTPS (HTTP Secure) certificate used by the Delphix Virtualization Engine. There are two methods of replacing the certificate. The key difference between the two is whether Delphix or the user is providing the key pair (public and private key).

Delphix Provided Key Pair

Use the following instructions to provide an HTTPS certificate chain for a key pair created by the Delphix Engine. Once the key pair is created users can download a Certificate Signing Request (CSR) to generate a signed certificate from the CA of their choice. This is done using the "create" operation in the "/service/tls/csr" API as seen below when using the CLI.

```
hostname.domainname> service tls csr
hostname.domainname service tls csr> create
hostname.domainname service tls csr create *> ls
Properties
  type: CertificateSigningRequestCreateParameters
  dname:
    type: X500DistinguishedNameComposite
    dname: (required)
  :
    type: EndEntityHttps
  forceReplace: false
  keyPair:
    type: RsaKeyPair
    keySize: 2048
    signatureAlgorithm: SHA256withRSA
```

The first key property is the dname. This will be used as the subject name of the CSR and resulting X.509 certificate unless it is changed when the certificate is signed. Delphix supports two different formats for dname:

- a composite string
- a list of fields

Use the composite string as follows:

```
hostname.domainname service tls csr create *> set dname.dname="CN=Delphix CA,
O=Delphix, C=US"
hostname.domainname service tls csr create *> ls
Properties
  type: CertificateSigningRequestCreateParameters
  dname:
    type: X500DistinguishedNameComposite (*)
    dname: CN=Delphix CA, O=Delphix, C=US (*)
```

Use the list of field formats as follows:

```
hostname.domainname service tls csr create *> set
dname.type=X500DistinguishedNameFields
hostname.domainname service tls csr create *> ls
Properties
  type: CertificateSigningRequestCreateParameters
  dname:
    type: X500DistinguishedNameFields (*)
    city: (unset)
    commonName: Delphix CA (*)
    country: US (*)
    organization: Delphix (*)
    organizationUnit: (unset)
    stateRegion: (unset)
```

The only required field is the commonName (CN).

The only currently supported type for endEntity is EndEntityHttps.

The next property is forceReplace. By default, this is false and means Delphix will not replace the active key pair and certificate with the newly generated keypair and self-signed certificate. If the user wants to replace the active key pair right away before the signed certificate has been created this can be set to **true**.

The final property keyPair impacts the generated key pair. When creating a new key pair the engine supports two algorithms:

- **RSA** - The supported signature algorithms are SHA256withRSA, SHA384withRSA, and SHA512withRSA . The valid key sizes range from 2048 to 4096.
- **ECDSA** - The supported signature algorithms are SHA256withECDSA, SHA384withECDSA, and SHA512withECDSA. The valid key sizes range from 256 to 571

Once the create operation has completed you can get the CSR in PEM format by selecting the CSR object and looking at the requestInPem property:

```
requestInPem: -----BEGIN CERTIFICATE REQUEST-----
MIIBezCB0gIBADAhMR8wHQYDVQDEZXiYmFrZXIuZGMYLmRlbHBoaXguY29tMIGN
MBAGByqGSM49AgEGBSuBBAAmA4GSAAQBU5WY9+GkCTFvbGHTNJDb/QM3t4YI/9S6
fhCJELx7SbJNti2n0l3mCePenyUuBY9m6BWvUQzlhawZG5YAJ9WdcM+IIPciiNsD
Xmw0eFH05z6yTLMnfBYYZKFbpu/dcK5V8WoItrIC7jTvg/k6jf/WeD+dmyIMQ0Z7
VmwnD6RsaAs7T5lajXkurwPfq5MnsmgADAKBggqhkJOPQDAwOBLwAwgZMCSAGM
quqcnIAXIRDxQ+BzzSywNtozn5ihtfFxtTF/EW/ARBib2l9hq0pwHrIinnLvjW9u
avpAH1pkWHx1w0/06W6DCZAPIIL3ugJHKsScJqsvaeZzVqJVfQt8g42cL9hKc7ic
HLhuAyMGQOXrEdLbOxtOH6SiExnyEv2Y9LHHYYgRafgGz0oA5tx+mrkr9J+zm8Y=
-----END CERTIFICATE REQUEST-----
```

Once the CSR has been signed and turned into an X.509 Certificate you can replace the certificate using the "service/tls/endEntityCertificate" API. To replace using the CSR method begin by setting the correct type of replace parameters as seen below:

```
hostname.domainname service tls endEntityCertificate> replace
hostname.domainname service tls endEntityCertificate replace *> set
type=EndEntityCertificateReplaceChainParameters
```

```
hostname.domainname service tls endEntityCertificate replace *> ls
Properties
  type: EndEntityCertificateReplaceChainParameters
  chain:
    type: PemCertificateChain
    chain: (required)
  endEntity:
    type: EndEntityHttps
```

The "chain" property must contain a list of the entire trust chain from the newly generated end-entity certificate to the root CA.

The CLI might not always interpret newline characters in PEM certificates correctly. Therefore, it is highly recommended to find and replace all newlines ('\n') with an empty string ("") prior to pasting the PEM certificate into the CLI.

To do this in the CLI first run:

```
hostname.domainname service tls endEntityCertificate replace *> edit chain.chain
Then `add` and `set contents` to the PEM certificate for each certificate in the
chain.
hostname.domainname service tls endEntityCertificate replace chain.chain *> add
```

When adding multiple certificates, use the command back after each add. After the final add, enter back and then commit.

The order that the PEM certificates are added to the list does not matter.

Finally, once the certificate chain has been successfully replaced you will need to **restart** the Delphix Engine.

Customer Provided Key Pair

This section describes the steps to take if you are replacing the HTTPS with your own key pair and certificate.

1. To start, you need to add the key pair and full certificate chain as an entry in a file in JKS or PKCS #12 format.
2. Then, send a file upload request to the following endpoint:

```
hostname.domainname service tls endEntityCertificate requestKeyPairAndCertChainUpload
*> ls
Properties
  type: CertificateUploadParameters
  alias: alias_in_keystore (*)
  keypass: (unset)
  keystoreType: JKS
  storepass: ***** (*)
hostname.domainname service tls endEntityCertificate requestKeyPairAndCertChainUpload
*> commit
type: FileUploadResult
token: 8f4361c5-019c-4fee-9306-b7c85e977cf4
url: /resources/json/delphix/data/upload
```

The **alias** field is where the key pair and certificate is saved in your JKS or PKCS #12 store.

The **keypass** field is the password for the given alias' key. If not set, it uses the keystore's password.

The **storepass** field is the keystore's password.

3. Then, establish a session from the host with the keystore to the Delphix Engine.

Choose the location of the cookies, and determine the API version (command example uses 1.9.2):

```
curl -c <path/to/cookies> -X POST --data '{ "type": "APISession", "version":
{ "type": "APIVersion", "major": 1, "minor": 9, "micro": 2 } }' -H "Content-Type:
application/json" http://<delphix_engine_url>/resources/json/delphix/session
```

4. Login to the Delphix Engine using the established session as a domain or system admin:

```
curl -b <path/to/cookies> -X POST --data '{ "type": "LoginRequest", "username":
"sysadmin", "password": "sysadmin" }' -H "Content-Type: application/json" http://
<delphix_engine_url>/resources/json/delphix/login
```

5. Send the file upload request with the location of your keystore and token from above:

```
curl -b <path/to/cookies> -X POST -F "file=@<path/to/keystore>" -F "token=8f4361c5-01
9c-4fee-9306-b7c85e977cf4" http://<delphix_engine_url>/resources/json/delphix/data/
upload
```

6. You can now replace the HTTPS end-entity certificate with the keystore you have uploaded, identified by the token:

```
hostname.domainname service tls endEntityCertificate replace *> set
type=EndEntityCertificateReplaceKeystoreParameters
hostname.domainname service tls endEntityCertificate replace *> set
token=8f4361c5-019c-4fee-9306-b7c85e977cf4
hostname.domainname service tls endEntityCertificate replace *> ls
Properties
  type: EndEntityCertificateReplaceKeystoreParameters (*)
  endEntity:
    type: EndEntityHttps
    token: b0e889ff-847a-4d7d-bd17-c1292ddb63e (*)
hostname.domainname service tls endEntityCertificate replace *> commit
```

7. Once the key pair and certificate chain has been successfully replaced you will need to **restart** the Delphix Engine for the changes to take effect.

CLI Cookbook: Configuring Key-Based SSH Authentication for Automation

This topic describes how to use CLI commands to configure individual users with SSH keys to allow for password-less authentication from a remote host to the CLI in an automated environment.

What is SSH Key-based Authentication?

Secure Shell (SSH) is a connection method used to log into UNIX or Linux servers remotely. With Delphix, it is used to connect to the Delphix Command Line Interface (CLI) from a remote computer. This normally requires a password on each connection; however, it is possible to use Key-based Authentication to avoid the password requirement and allow the automation of Delphix commands.

Key-based Authentication relies on a public/private key pair generated on the client system. The private key allows access to any server acknowledging the matching public key as being authorized to login. In order to configure this, a public/private key pair must be created, and the resulting public key should be added to the Delphix server using the CLI.

 The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Procedure

1. Consult your client documentation for information on generating a public/private key pair. The `ssh-keygen` program is typical on UNIX platforms. If you need details on `ssh-keygen` usage or have unique requirements (such as named RSA keys), see [Third-Party SSH Key Generation Example](#). If you already have a public/private key pair generated on your system, you can skip to step 2.
2. Connect as the user you wish to configure or as a Delphix administrator.

Connecting to Namespaces

When you connect to the Delphix Engine with the CLI, you should specify the appropriate namespace (either `DOMAIN` or `SYSTEM`). See [Connecting to the CLI](#) for more information.

3. Select the current user, or select a specific user if configuring another user as an administrator.

```
delphix> user current
```

4. Update the user and set the SSH key.

```
delphix user "admin"> update
delphix user "admin" update *> set publicKey="[PASTE KEY]"
delphix user "admin" update *> commit
delphix>
```

Avoid Newline Characters with Public Key Entry

The public key value, which can be quite long, must be entered as a single string with no newlines. When copying and pasting the public key, be sure to avoid introducing any newline characters. For more information on how to manage multiple public keys for password-less user authentication on Delphix, please visit this [Knowledge Base](#) article.

5. Verify you can authenticate through the Delphix CLI without a passphrase.

Example Using Default SSH Key

```
ssh admin@DOMAIN@delphix-server.example.com
Last login: Thu Dec 13 22:16:28 2012 from 192.168.0.2
delphix>
```

Example Using a Non-default SSH Key File Located at path/to/delphix_key

```
ssh -i path/to/delphix_key admin@DOMAIN@delphix-server.example.com
Last login: Thu Dec 13 22:16:28 2012 from 192.168.0.2
delphix>
```

Third Party SSH Key Generation Example

This topic provides examples of SSH key generation for password-less login to the Delphix command-line interface (CLI).

- [Prerequisites](#)
- [Procedure](#)
 - [Using an Existing SSH Key With No Passphrase](#)
 - [Generating a New SSH Key Stored in the Default Location](#)
 - [Generating a New SSH Key Stored in a Non-Default Location](#)

Prerequisites

These operations are performed as a command-line user on a non-Delphix host, where SSH is installed. In the remainder of the document we will use the terms "username" and "hostname" in the following way:

- `username` – the existing command-line user on the non-Delphix host
- `host name` – the existing non-Delphix host

The examples below should work with a variety of SSH distributions. However, your distribution may behave differently. If you are unable to follow these instructions successfully, consult with your system administrator and/or your operating system or SSH client vendor. Please note that only RSA protocol 2 keys are supported.

Procedure

This example details three ways to set up password-less authentication:

- Using an existing SSH key with no passphrase
- Generating a new SSH key stored in the default location
- Generating a new SSH key stored in a non-default location

Using an Existing SSH Key With No Passphrase

- ✓ To use an existing SSH key for password-less authentication, that key must have been established with no passphrase. If the existing SSH key has a passphrase associated with it, instead follow the instructions for generating a new SSH key stored in a non-default location below.

1. Login as `username` to `hostname`.
2. Within the `~username/.ssh/` directory, there will be one or more files whose names end in the `.pub` extension, for example:
 - a. `id_dsa.pub`
 - b. `id_rsa.pub`
 - c. `identity.pub`

The differing names correspond to different types of keys SSH may use. Note that from version 5.0 of Delphix, only RSA protocol 2 keys are supported.

3. Input the contents of the `id_rsa .pub` file into the `publicKey` field as described in [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#).

Generating a New SSH Key Stored in the Default Location

1. Login as `username` to `hostname`.

2. Verify that there are no files with the `.pub` extension in the `~username/.ssh/` directory. If there are, you must either use the existing SSH key as described above or generate a new SSH key in a non-default location as described below.
3. Create your RSA protocol 2 keys with no passphrase.
 - a. Run `ssh-keygen -t rsa`
 This command creates two new files in `~username/.ssh/`: `id_rsa` and `id_rsa.pub`
4. You will be prompted to enter a **passphrase** and to confirm the passphrase. Simply hit **Enter** twice to create a key with no passphrase for password-less authentication.
5. Input the contents of the created `id_rsa .pub` file in the `~username/.ssh/` directory into the `publicKey` field as described in [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#).

Generating a New SSH Key Stored in a Non-Default Location

1. Login as `username` to `hostname`.
2. Choose the filename for your new SSH key files, for example: `delphix_key` and `delphix_key.pub`
3. Create your keys with no passphrase.
 - a. Run `ssh-keygen -t rsa -f delphix_key`
 This command creates two new files `delphix_key` and `delphix_key.pub`
4. You will be prompted to enter a **passphrase**, and to confirm the passphrase. Simply hit **Enter** twice to create a key with no passphrase for password-less authentication.
5. Input the contents of the created `.pub` file in the `~username/.ssh/` directory into the `publicKey` field as described in [CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#).
6. On subsequent login attempts to the Delphix CLI, specify the private SSH key to be used with the `-i` argument, for example:

```
ssh -i ~/delphix_key username@delphixserver.company.com
```

CLI Cookbook: Setting Up SSH Key Authentication for UNIX Environment Users

This topic describes adding public-key authentication for a UNIX environment user, thus allowing the Delphix server to connect to your UNIX Environments without an explicit password. This method uses the Delphix CLI in order to set up the environment user and gather SSH keys.

UNIX host environments (and Oracle cluster environments) can have users configured to use SSH-key-based authentication instead of the traditional password authentication method.

Prerequisites

- You must be able to log into the remote host (or all hosts of an Oracle cluster) and have write access to the `~/.ssh/authorized_keys` file within the desired user's home directory.

Option 1: System Key

Within Delphix, there is a per-system SSH public key that can be placed into the `~/.ssh/authorized_keys` file of the remote user. Once this has been done, the Delphix environment user can be configured to use the private key instead of an explicit password. Note that it is also possible to configure an environment to use this system key in the Delphix Management application by navigating to **Manage > Environments** and selecting **Public Key** as the **Login Type** for the environment. For details, see [Managing Environments](#).

1. Get the current system public key:

```
delphix> system get sshPublicKey
ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAQEase1M7uJX44lVPBljhnxB6MZUTx8VF6cupaVATg120lQonIqx29l
P+Mwp0AWh7C983IDoYDo+AY7RXpcFP9nKksiJnGSGiK6wo9RIiqSnF1x/
VXNkTt2/67RVofoiui4W5fuxD4h0IvoTr47Bg1hh9L6nhP0tnUvS/
rusHFJ+ogxGHm46mwNlgUJUGmLTNao+W0YU693HRLukEch01t4k6oLVGaC0eLjYlgBf0Z5XiIcBX6ZW
qVHAhwMinVjAvmfQhirAgCI7gYrd5/PwNl/DC8xyhWuxd2jgA7sSPeRqWY0JHt/
xcmdpIaPxTwtxQLKTnPxrFrQd+l4uf6LKxr5g7w== root@delphix
```

2. Add this key (starting with `ssh-rsa`) to the remote user's `~/.ssh/authorized_keys` file. You will need to get access to this file using an alternate authentication mechanism (such as logging in as the user with a password or logging in as an administrator). Depending on the target OS, you may need to do the following:
 - a. If the directory does not exist:

```
$ mkdir ~/.ssh
```

- b. If creating the file or directory as an administrator:

```
# chown -R <username> <home>/.ssh
```

- c. If required by the host SSH configuration, ensure the directory is not world-readable:

```
$ chmod 600 ~/.ssh/authorized_keys
$ chmod 755 ~
```

3. Create or edit an environment user:

```
delphix> environment user create
```

4. Set the user environment and name:

```
delphix environment user create *> set environment=environment1
delphix environment user create *> set name=username
```

5. Set the user credential type to `SystemKeyCredential` :

```
delphix environment user create *> set credential.type=SystemKeyCredential
```

6. Commit the results:

```
delphix environment user create *> commit
```

Option 2: Per-environment Key Pair

Each environment user can also be configured to use an SSH key pair provided via the CLI or API.

1. Add the public key to the remote user's `~/.ssh/authorized_keys` file. You will need to get access to this file using an alternate authentication mechanism (such as logging in as the user with a password or logging in as an administrator). Depending on the target OS, you may need to do the following:
 - a. If the directory does not exist:

```
$ mkdir ~/.ssh
```

- b. If creating the file or directory as an administrator:

```
# chown -R <username> <home>/.ssh
```

- c. If required by the host SSH configuration, ensure the directory is not world-readable:

```
$ chmod 600 ~/.ssh/authorized_keys
$ chmod 755 ~
```

2. Create or edit an environment user:

```
delphix> environment user create
```

3. Set the user environment and name:

```
delphix environment user create *> set environment=environment1
delphix environment user create *> set name=username
```

4. Set the user credential type to `KeyPairCredential` :

```
delphix environment user create *> set credential.type=KeyPairCredential
```

5. Set the private and public keys:

```
delphix environment user create *> set credential.privateKey="-----BEGIN ..."  
delphix environment user create *> set credential.publicKey="ssh-rsa AA..."
```

(these example values were trimmed for brevity)

6. Commit the results:

```
delphix environment user create *> commit
```

CLI Cookbook: Configuring SSH Host Verification for UNIX Environments

This topic describes how to configure SSH host verification when authenticating to UNIX environments, which lets the Delphix server ensure it connects to the intended environment hosts. This method uses the Delphix CLI to set the SSH key or fingerprint of each host. Currently, it is only possible to configure SSH host verification via the CLI or the Web Service API.

When an SSH key or fingerprint is specified for an environment host, the Delphix server will use it when connecting to that host to verify that host. If the key or fingerprint does not match the information presented by that host, the Delphix server will close that connection and report the problem to the user.

The key types supported by the Delphix server are `RSA` (`ssh-rsa`), `DSA` (`ssh-dsa`), `ECDSA` (`ecdsa-sha2-nistp256`) and `ED25519` (`ssh-ed25519`). The fingerprint types supported are `SHA256` and `SHA512` ; the `MD5` type is considered insecure and, therefore, is not supported.

Prerequisites

- To obtain the SSH public key or fingerprint of a host remotely from another machine, you must have the `ssh-keyscan` and `ssh-keygen` utilities.
- To obtain the SSH public key or fingerprint directly from a host, you must be able to log into that host.

Obtaining an SSH key or fingerprint

1. Remotely: List the SSH public keys of the host using the standard utility `ssh-keyscan` and choose one of them. For example:

```
$ ssh-keyscan example.environment.host
# example.environment.host:22 SSH-2.0-OpenSSH_7.4
example.environment.host ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDBsOAAokSR067jI28syRmX0wY/fKIboLLu/
ofk6BzYLKtkMaK1QC78/6QleLIJUP5HdK8E7Um/iM1JMxry4h9Rl13onY0uJVZkDB9wnJiztSu/
Wl9Eqbt59TU1vGmp/4uLWS3PISl7bxs+l43HzsrjM4dTs2efQ7sLWoW86CDlL7Je4va65/
aopvifxKZeZkT0srB3L8VzHKw9+NJOumy1CI3DIBiICURJd4WZ10IH5TFUDRaUFAC/trzW1gvJY/
Whp892tPHeKyP32h0ZNIc7oDPx2boZauJVR6/
BHmKpmLlhkPpEqfZP8JW+JNsNnLr9BEmwJXaEpwnua1BUi8F
...
```

where the key is the Base64-code string to the right of the key type. In this example, the RSA SSH public key is the string starting in "AAAAB3Nza" and ending in "a1BUi8F".

2. Alternatively, from the host: Log into the host and print the file contents of your public key of choice. For example:

```
$ cat /etc/ssh/ssh_host_rsa_key.pub
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDBsOAAokSR067jI28syRmX0wY/fKIboLLu/
ofk6BzYLKtkMaK1QC78/6QleLIJUP5HdK8E7Um/iM1JMxry4h9Rl13onY0uJVZkDB9wnJiztSu/
Wl9Eqbt59TU1vGmp/4uLWS3PISl7bxs+l43HzsrjM4dTs2efQ7sLWoW86CDlL7Je4va65/
aopvifxKZeZkT0srB3L8VzHKw9+NJOumy1CI3DIBiICURJd4WZ10IH5TFUDRaUFAC/trzW1gvJY/
Whp892tPHeKyP32h0ZNIc7oDPx2boZauJVR6/
BHmKpmLlhkPpEqfZP8JW+JNsNnLr9BEmwJXaEpwnua1BUi8F
```

3. If a fingerprint is preferred, use `ssh-keygen` in conjunction with the above commands. For example, remotely:

```
$ ssh-keyscan example.environment.host | ssh-keygen -E sha256 -lf -
...
2048 SHA256:8Cx8cBg/pSbkId3uu2vATEugkAXcm+Ruu9hu660XEGI
example.environment.host (RSA)
...
```

where the fingerprint is " `SHA256:8Cx8cBg/pSbkId3uu2vATEugkAXcm+Ruu9hu660XEGI` " (the string between the key size and the hostname).

Alternatively, from the host:

```
$ cat /etc/ssh/ssh_host_rsa_key.pub | ssh-keygen -E sha256 -lf -
2048 SHA256:8Cx8cBg/pSbkId3uu2vATEugkAXcm+Ruu9hu660XEGI user@environment.host
(RSA)
```

Configuring SSH host verification during environment creation

The default SSH verification strategy is `SshAcceptAlways`, which always trusts the key or fingerprint presented by a remote host. The procedure to change this strategy to perform fingerprint-based host verification for single-host Unix environments is:

1. Set the new strategy to `SshVerifyFingerprint`:

```
delphix environment create * > edit hostParameters.host.sshVerificationStrategy
delphix environment create hostParameters.host.sshVerificationStrategy * > set
type=SshVerifyFingerprint
```

2. Set the key type and fingerprint type. For example:

```
delphix environment create hostParameters.host.sshVerificationStrategy * > set
keyType=RSA
delphix environment create hostParameters.host.sshVerificationStrategy * > set
fingerprintType=SHA256
```

3. Set the fingerprint. For example:

```
delphix environment create hostParameters.host.sshVerificationStrategy * > set
fingerprint=SHA256:8Cx8cBg/pSbkId3uu2vATEugkAXcm+Ruu9hu660XEGI
```

4. Alternatively, you can specify the key itself using the `SshVerifyRawKey` strategy. For example:

```
delphix environment create hostParameters.host.sshVerificationStrategy * > set
type=SshVerifyRawKey
```

```
delphix environment create hostParameters.host.sshVerificationStrategy *> set
keyType=RSA
delphix environment create hostParameters.host.sshVerificationStrategy *> set
rawKey=AAAAB3NzaC1yc2EAAAADAQABAAQDBs0AAokSR067jI28syRmX0wY/fKIboLLu/
ofk6BzYLKtkMaK1QC78/6QleLIJUP5HdK8E7Um/iM1JMxry4h9Rl13onY0uJVZkDB9wnJiztSu/
Wl9Eqbt59TU1vGmp/4uLWS3PISl7bxs+l43HzsrjM4dTsfQ7sLWoW86CDLL7Je4va65/
aopvifxKZeZkT0srB3L8VzHKw9+NJOumy1CI3DIBiICURJd4WZ10IH5TFUDRaUFac/trzW1gvJY/
Whp892tPHekyP32h0ZNIc7oDPx2boZauJVR6/
BHmKpmLlhkPpEqfZP8JW+JNsNnLr9BEmwJXaEpwnua1BUii8F
```

- When you are done specifying all other environment creation parameters, create the environment:

```
delphix environment create *> commit
```

- If you are creating a Unix cluster, the procedure to start editing the SSH verification settings for the first node in that cluster is similar to the single-host case. For example:

```
delphix> environment create
delphix environment create *> set type=OracleClusterCreateParameters
delphix environment create *> edit nodes
delphix environment create nodes *> add
delphix environment create nodes 0 *> edit
delphix environment create nodes 0 *> edit
hostParameters.host.sshVerificationStrategy
# configure SSH verification settings
```

Note that only one node (host) can be specified and configured when creating a Unix cluster environment. The SSH verification settings for the remaining hosts can only be specified afterward by editing them via "`host select <hostname> update`", once the corresponding nodes have been discovered or added. See the next section.

Configuring SSH host verification for existing hosts

For any Unix environment host, whether it is single or part of a cluster, you can set up or change its configuration for SSH verification after the environment has been added by editing the host. For example:

```
delphix> host select example.environment.host update sshVerificationStrategy
delphix host 'example.environment.host' update sshVerificationStrategy *> edit
sshVerificationStrategy
# configure SSH verification settings
```

Testing SSH host verification

It is possible to configure an SSH key or fingerprint when performing a connectivity test to a Unix host. This can be done without even creating an environment for that host. For example:

```
delphix> connectivity ssh
# configure address and credentials
...
```

```
delphix connectivity ssh *> edit sshVerificationStrategy
# configure SSH verification settings
```

SSH host verification errors

When the Delphix server initiates an SSH connection to a host, if SSH host verification is configured (i.e. the verification strategy is not the default `SshAcceptAlways`), the server will first check the key presented by the host. Only if this check passes, the server will attempt to authenticate. Therefore, a host key verification failure will be reported as "Unrecognized key or fingerprint" to the user before any authentication failure. For example:

```
delphix connectivity ssh *> set credentials.password=<BAD PASSWORD>
delphix connectivity ssh *> set sshVerificationStrategy.rawKey=<BAD KEY>
delphix connectivity ssh *> commit
  Error: Unrecognized key or fingerprint returned by host "example.environment.host".
  Action: Contact your administrator as this could be a man-in-the-middle attack.
  Otherwise, correct the host's ssh key/fingerprint configuration.
delphix connectivity ssh *> set sshVerificationStrategy.rawKey=<GOOD KEY>
delphix connectivity ssh *> commit
  Error: Unable to authenticate to host "example.environment.host" using username
  "myusername".
  Action: Check the username and password and try again.
```

CLI Cookbook: Password Policies

These topics describe command-line interface procedures for managing password policies.

- [Changing the Current Password Policy to an Available Password Policy](#)
- [Creating Custom Password Policies](#)
- [Deleting a Custom Password Policy](#)
- [Updating a Custom Password Policies](#)
- [Updating a Password](#)
- [Viewing Available Password Policies](#)
- [Viewing the Current Password Policy Requirements](#)

Changing the Current Password Policy to an Available Password Policy

This topic describes how to change the current password policy to any of the available password policies.

Procedure

1. Log in to the Delphix Engine as the administrator user and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Run the changeCurrentPasswordPolicy command.

```
delphix passwordPolicy> changeCurrentPasswordPolicy  
delphix passwordPolicy changeCurrentPasswordPolicy *>
```

3. Set the required parameters.

```
delphix passwordPolicy changeCurrentPasswordPolicy *> set userType=DOMAIN  
delphix passwordPolicy changeCurrentPasswordPolicy *> set  
currentPasswordPolicy=MyCustomPasswordPolicy
```

4. Commit the operation.

```
delphix passwordPolicy changeCurrentPasswordPolicy *> commit
```

Creating Custom Password Policies

This topic describes how to create custom password policies.

Procedure

1. Log in to the Delphix Engine as the administrator user and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Run the create command.

```
delphix passwordPolicy> create  
delphix passwordPolicy create *>
```

3. Set the required parameters.

```
delphix passwordPolicy create *> set name=MyCustomPasswordPolicy minLength=5  
reuseDisallowLimit=1 uppercaseLetter=true lowercaseLetter=false digit=true  
symbol=true disallowUsernameAsPassword=true
```

4. Commit the operation.

```
delphix passwordPolicy create *> commit
```

Deleting a Custom Password Policy

This topic describes how to delete custom password policies.

Procedure

1. Log in to the Delphix Engine as the administrator user and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Select the custom password policy to be deleted.

```
delphix passwordPolicy> select MyCustomPasswordPolicy  
delphix passwordPolicy 'MyCustomPasswordPolicy'>
```

3. Run the delete command.

```
delphix passwordPolicy 'MyCustomPasswordPolicy'> delete
```

4. Commit the operation.

```
delphix passwordPolicy 'MyCustomPasswordPolicy' delete *> commit
```

Updating a Custom Password Policies

This topic describes how to update custom password policies.

Procedure

1. Log in to the Delphix Engine as the administrator user and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Select the custom password policy to be updated.

```
delphix passwordPolicy> select MyCustomPasswordPolicy
```

3. Run the update command.

```
delphix passwordPolicy 'MyCustomPasswordPolicy'> update
```

4. Set the required parameters to new values.

```
delphix passwordPolicy 'MyCustomPasswordPolicy' update *> set minLength=8
```

5. Commit the operation

```
delphix passwordPolicy 'MyCustomPasswordPolicy' update *> commit
```

Updating a Password

This topic describes how to update the password.

Procedure

1. Log in to the Delphix Engine and switch to the user context.

```
delphix> user
```

2. Run the 'current' command to select the currently logged in user.

```
delphix user> current
```

3. Run the updateCredential command.

```
delphix user 'xyz'> updateCredential
```

4. Set the required credentials.

```
delphix user 'xyz' updateCredential *> set oldCredential.password=XXXXXXXXX  
delphix user 'xyz' updateCredential *> set newCredential.password=XXXXXXXXX
```

5. Commit the operation.

```
delphix user 'xyz' updateCredential *> commit
```

Viewing Available Password Policies

This topic describes how to view available password policies

Procedure

1. Log in to the Delphix Engine as the administrator user and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Run the list command.

```
delphix passwordPolicy> list
REFERENCE          NAME  MINLENGTH  REUSEDISALLOWLIMIT  UPPERCASELETTER
LOWERCASELETTER  DIGIT  SYMBOL  DISALLOWUSERNAMEASPASSWORD

PASSWORDPOLICY-0  NONE  1          0                    false           false
      false false false

delphix> passwordPolicy
```

Viewing the Current Password Policy Requirements

This topic describes how to view current password policy requirements.

Procedure

1. Log in to the Delphix Engine and switch to the passwordPolicy context.

```
delphix> passwordPolicy
```

2. Run the currentRequirements command.

```
delphix passwordPolicy> currentRequirements
```

3. Commit the operation (following is an example, not a suggestion).

```
delphix passwordPolicy currentRequirements *> commit
Password must:
- Contain at least 8 characters.
- Differ from the last password.
- Not contain the user name or reverse user name.
- Contain at least one uppercase letter.
- Contain at least one numeric digit.
- Contain at least one symbol.
```

CLI Cookbook: Changing the Default admin username

⚠ The default domain user created on Delphix Engines is now admin instead of delphix_admin. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

To change the default admin username.

1. ssh into your engine
2. Go to users and select admin.
3. Change the name to admin.
4. Commit the operation.

Example

```
delphix.engine> user select admin

delphix.engine user 'admin'> update

delphix.engine user 'admin' update *> set name=dlpxadmin

delphix.engine user 'admin' update *> commit

delphix.engine user 'dlpxadmin'>
```

CLI Cookbook: Disabling an admin User

⚠ The default domain user created on Delphix Engines is now admin instead of delphix_admin. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling delphix_admin.

To disable your default admin.

1. ssh into your engine
2. Go to users and select admin.
3. Select the default admin and use 'disable' command
4. Commit the operation.

```
delphix.engine> user select admin
delphix.engine user 'admin'> disable
delphix.engine user 'admin' disable *> commit
delphix.engine user 'admin'>
```

CLI Cookbook: System Administration

These topics describe various system administration tasks that can be performed with the command-line interface, such as changing the name of the <default> group and setting up network connectivity.

This section covers the following topics:

- [CLI Cookbook: Configuring A Second Network Interface](#)
- [Removing IP Addresses from a Network Interface](#)
- [CLI Cookbook: Adding A Static Route](#)
- [CLI Cookbook: Changing the Default Group Name](#)
- [CLI Cookbook: How to Change a Delphix User Password](#)
- [CLI Cookbook: Retrieve Capacity Information](#)
- [CLI Cookbook: View Storage Test Results](#)
- [CLI Cookbook: How to Change IP Address of Delphix Engine](#)
- [CLI Cookbook: About Alert Notifications](#)
- [CLI Cookbook: Obtaining CPU Performance Information Using CLI](#)
- [CLI Cookbook: Rebooting the Delphix Engine via CLI](#)
- [CLI Cookbook: Disabling User-click Analytics](#)
- [CLI Cookbook: Changing the API Version](#)

CLI Cookbook: Configuring A Second Network Interface

This topic describes how to configure a static IP address on a second network interface.

Procedure

1. Add a NIC to the Delphix virtual machine. The specific procedure will depend on the platform. For example, on VMware, a VMXNET3 virtual network adapter can be added dynamically from vSphere or other administrative interfaces, and the Delphix Engine will recognize the new NIC without a reboot. On other platforms, a reboot may be required for the Delphix Engine to recognize the new virtual hardware.
2. Log in to the Delphix Engine as the sysadmin user and switch to the network interface context. Then use the `list` command to view the available network interfaces, and select the new interface to be configured.

```
delphix network interface> list
NAME
vmxnet3s0
vmxnet3s1
delphix network interface> select vmxnet3s1
fender.talisker network interface "vmxnet3s1"> get
  type: NetworkInterface
  name: vmxnet3s1
  addresses: (empty)
  device: vmxnet3s1
  macAddress: 0:c:29:e5:4c:c1
  mtu: 1500
  mtuRange: 60-9000
  reference: NETWORK_INTERFACE-vmxnet3s1
  state: DOWN
```

3. Run the `update` command and configure a static address.

```
delphix network interface "vmxnet3s1"> update
delphix network interface "vmxnet3s1" update *> edit addresses.0
delphix network interface "vmxnet3s1" update addresses.0 *> set address=10.1.2.3/24
delphix network interface "vmxnet3s1" update addresses.0 *> get
  type: InterfaceAddress (*)
  address: 10.1.2.3/24 (*)
  addressType: STATIC (*)
```

4. `Commit` the operation.

```
delphix network interface "vmxnet3s1" update addresses.0 *> commit
delphix network interface "vmxnet3s1"> get
  type: NetworkInterface
  name: vmxnet3s1
  addresses:
    0:
```

```
type: InterfaceAddress
address: 10.1.2.3/24
addressType: STATIC
state: OK
device: vmxnet3s1
macAddress: 0:c:29:e5:4c:c1
mtu: 1500
mtuRange: 60-9000
reference: NETWORK_INTERFACE-vmxnet3s1
state: OK
```

Removing IP Addresses from a Network Interface

This topic describes how to remove one or more IP addresses from a network interface.

Removing all IP addresses

Procedure

1. Log in to the Delphix Engine as the sysadmin user and switch to the network interface context. Then use the `list` command to view the available network interfaces, and select the interface.

```
delphix> network interface
delphix network interface> list
NAME      MACADDRESS      MTU
ens160    02:dc:02:00:7f:06 1500
delphix network interface> select ens160
delphix network interface 'ens160'> get
  type: NetworkInterface
  name: ens160
  addresses:
    0:
      type: InterfaceAddress
      address: 10.43.70.135/16
      addressType: DHCP
      enableSSH: true
      sessionInUse: false
      state: OK
  device: ens160
  macAddress: 02:dc:02:00:7f:06
  mtu: 1500
  mtuRange: 68-9000
  reference: NETWORK_INTERFACE-ens160
  state: OK
```

2. Run the `update` command and unset the addresses property to remove all addresses. Note that in this example, we only had one address to remove, but this procedure removes all addresses even if more than one were configured.

```
delphix network interface 'ens160'> update
delphix network interface 'ens160' update *> unset addresses
delphix network interface 'ens160' update *> get
  type: NetworkInterface
  name: ens160
  addresses: (unset) (*)
  mtu: 1500
```

3. `Commit` the operation.

```
delphix network interface 'ens160' update *> commit
```

```
delphix network interface 'ens160'> get
type: NetworkInterface
name: ens160
addresses: (empty)
device: ens160
macAddress: 02:dc:02:00:7f:06
mtu: 1500
mtuRange: 68-9000
reference: NETWORK_INTERFACE-ens160
state: OK
```

Removing a specific IP address

1. Log in to the Delphix Engine as the sysadmin user and switch to the network interface context. Then use the `list` command to view the available network interfaces, and select the interface. Note that in this example, there are two addresses, a static address, and a DHCP address.

```
delphix> network interface
delphix network interface> list
NAME      MACADDRESS      MTU
ens160    02:dc:02:00:7f:06 1500
delphix network interface> select ens160
delphix network interface 'ens160'> get
type: NetworkInterface
name: ens160
addresses:
  0:
    type: InterfaceAddress
    address: 10.11.12.13/24
    addressType: STATIC
    enableSSH: true
    sessionInUse: false
    state: OK
  1:
    type: InterfaceAddress
    address: 10.43.70.135/16
    addressType: DHCP
    enableSSH: true
    sessionInUse: false
    state: OK
device: ens160
macAddress: 02:dc:02:00:7f:06
mtu: 1500
mtuRange: 68-9000
reference: NETWORK_INTERFACE-ens160
state: OK
```

2. Run the `update` command and unset one of the addresses. In this example, only the static address is removed, retaining the DHCP address.

```

delphix network interface 'ens160'> update
delphix network interface 'ens160' update *> get
  type: NetworkInterface
  name: ens160
  addresses:
    0:
      type: InterfaceAddress
      address: 10.11.12.13/24
      addressType: STATIC
      enableSSH: true
    1:
      type: InterfaceAddress
      address: 10.43.70.135/16
      addressType: DHCP
      enableSSH: true
  mtu: 1500
delphix network interface 'ens160' update *> unset addresses.0
delphix network interface 'ens160' update *> get
  type: NetworkInterface
  name: ens160
  addresses:
    0:
      type: InterfaceAddress (*)
      address: 10.43.70.135/16 (*)
      addressType: DHCP (*)
      enableSSH: true (*)
  mtu: 1500

```

3. Commit the operation.

```

delphix network interface 'ens160' update *> commit
delphix network interface 'ens160'> get
  type: NetworkInterface
  name: ens160
  addresses:
    0:
      type: InterfaceAddress
      address: 10.43.70.135/16
      addressType: DHCP
      enableSSH: true
      sessionInUse: false
      state: OK
  device: ens160
  macAddress: 02:dc:02:00:7f:06
  mtu: 1500
  mtuRange: 68-9000
  reference: NETWORK_INTERFACE-ens160
  state: OK

```

CLI Cookbook: Adding A Static Route

This topic describes how to add a static route.

Procedure

1. Log in to the Delphix Engine as the sysadmin user and switch to the network route context.

```
delphix network route> list
DESTINATION      GATEWAY      OUTINTERFACE  PROTOCOL
default          172.16.0.1   ens192        DHCP
10.1.2.0/24      -            ens224        KERNEL
172.16.0.0/24    -            ens192        KERNEL
```

2. Run the `create` command to add a new route.

```
delphix network route> create
delphix network route create *> set destination=192.168.11.0/24
delphix network route create *> set gateway=10.1.2.1
delphix network route create *> get
  type: NetworkRoute
 destination: 192.168.11.0/24 (*)
 gateway: 10.1.2.1 (*)
 outInterface: (unset)
```

3. Optional outInterface Property

Setting the `outInterface` property is optional, as the system will automatically determine the output interface based on the gateway address provided, as shown below.

4. `Commit` the operation.

```
delphix network route create *> commit
delphix network route> list
DESTINATION      GATEWAY      OUTINTERFACE  PROTOCOL
default          172.16.0.1   ens192        DHCP
10.1.2.0/24      -            ens224        KERNEL
172.16.0.0/24    -            ens192        KERNEL
192.168.11.0/24  10.1.2.1     ens224        STATIC
```

CLI Cookbook: Changing the Default Group Name

This topic describes how to change the name of the default group <New Group> on the Delphix Engine as a simple example of CLI interactions. You must have delphix_admin credentials to perform this procedure.

Procedure

1. Switch to the group context and list groups on the system.

```
delphix> group
delphix group> list
NAME          DESCRIPTION
<New Group>  -
```

2. Select the default group and show current properties.

```
delphix group> select "<New Group>"
delphix group "<New Group>"> get
name: <New Group>
type: Group
description: (unset)
reference: GROUP-1
```

3. Run the `update` command and set the name.

```
delphix group "<New Group>"> update
delphix group "<New Group>" update *> set name=default
delphix group "<New Group>" update *> get
name: default (*)
type: Group
description: (unset)
reference: GROUP-1
```

4. `Commit` the operation.

```
delphix group "<New Group>" update *> commit
delphix group "default">
```

CLI Cookbook: How to Change a Delphix User Password

⚠ The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

1. ssh into your engine with a user that has Admin privileges.

```
ssh admin@delphix
```

2. Go to Users and select the User you would like to change the password of.

```
delphix > user
delphix user > ls
delphix user > select example_user
delphix user "example_user" > ls
```

3. Select `updateCredential` to allow you to change the password and set a new password.

```
delphix user "example user" > updateCredential
delphix user "example_user" updateCredential *> set newCredential.password=<new
password>
```

4. Commit the operation.

```
delphix user "example_user" updateCredential *> commit
```

Example:

```
ssh admin@delphixengine
delphixengine > user
delphixengine user > ls

Objects

NAME                EMAILADDRESS
-----
sysadmin            -
admin  no@delphix.com
test_user          no@delphix.com

Operations
```

```
create
```

```
current
```

```
delphixengine user > select test_user
```

```
delphixengine user "test_user" > ls
```

```
Properties
```

```
  type: User
```

```
  name: test_user
```

```
  authenticationType: NATIVE
```

```
  credential:
```

```
    type: PasswordCredential
```

```
    password: *****
```

```
  emailAddress: no@delphix.com
```

```
  enabled: true
```

```
  firstName: (unset)
```

```
  homePhoneNumber: (unset)
```

```
  isDefault: true
```

```
  lastName: (unset)
```

```
  locale: en_US
```

```
  mobilePhoneNumber: (unset)
```

```
  passwordUpdateRequested: false
```

```
  principal: test_user
```

```
  publicKey: (unset)
```

```
  reference: USER-2
```

```
  sessionTimeout: 30min
```

```
  userType: DOMAIN
```

```
  workPhoneNumber: (unset)
```

Operations

delete

update

disable

enable

updateCredential

```
delphixengine user "test_user" > updateCredential
```

```
delphixengine user "test_user" updateCredential *> set newCredential.password=<new  
password>
```

```
delphixengine user "test_user" update *> commit
```

CLI Cookbook: Retrieve Capacity Information

This topic describes how to gather capacity information from your Delphix Engine. This information includes:

- dSource Space Breakdown
- Virtual Object Space Breakdown
- Total Space

Procedure

1. Switch to the capacity system context.

```
delphix> capacity system
```

2. List the properties of this content.

```
delphix capacity system> ls
Properties
type: CurrentSystemCapacityData
source:
type: CapacityBreakdown
activeSpace: 940582400B
actualSpace: 1075381760B
descendantSpace: 134583808B
logSpace: 145920B
manualSpace: 0B
policySpace: 0B
syncSpace: 134583808B
timeflowUnvirtualizedSpace: 7725215744B
unvirtualizedSpace: 2624235520B
timestamp: 2015-12-11T11:49:18.998Z
totalSpace: 25568477184B
virtual:
type: CapacityBreakdown
activeSpace: 176684032B
actualSpace: 313768448B
descendantSpace: 0B
logSpace: 47820288B
manualSpace: 85958144B
policySpace: 0B
syncSpace: 85958144B
timeflowUnvirtualizedSpace: 5475587584B
unvirtualizedSpace: 2667149312B
```

For more information about capacity management in Delphix, visit [An Overview of Capacity and Performance Information](#).

CLI Cookbook: View Storage Test Results

1. Log into the CLI as the sysadmin.

```
ssh sysadmin@yourengine
```

2. Navigate to storage test.

```
delphix > storage  
delphix storage > test
```

3. List your tests and select the one that you would like to view and get the results.

```
delphix storage test > ls  
delphix storage test > select STORAGE_TEST-X  
delphix storage test STORAGE_TEST-X > result  
delphix storage test STORAGE_TEST-X result *> commit
```

CLI Cookbook: How to Change IP Address of Delphix Engine

This topic describes how to change an IP address on the Delphix Engine.

Procedure

1. Stop all running VDBs by clicking the **Stop** button on the VDB card.
2. Disable all dSources.
3. Log into the Delphix Engine as **sysadmin** user and switch to the network interface context. Then use the list command to view the available network interfaces, and select the public interface to be changed.

```
delphix> network
delphix network> setup
delphix network interface> list
NAME
vmxnet3s0
delphix network interface> select vmxnet3s0
delphix network interface 'vmxnet3s0'> get
  type: NetworkInterface
  name: vmxnet3s0
  addresses:
    0:
      type: InterfaceAddress
      address: 10.1.2.3/24
      addressType: STATIC
      enableSSH: true
      state: OK
  dataNode: DATA_NODE-34
  device: vmxnet3s0
  macAddress: 0:c:29:32:96:a3
  mtu: 1500
  mtuRange: 60-9000
  reference: NETWORK_INTERFACE-vmxnet3s0-DATA_NODE-34
  state: OK
```

4. Run the update command and update the address to the new IP address.

```
delphix network interface 'vmxnet3s0'> update
delphix network interface 'vmxnet3s0' update *> edit addresses.0
delphix network interface 'vmxnet3s0' update addresses.0 *> get
Properties
  type: InterfaceAddress
  address: 172.16.151.154/24
  addressType: STATIC
  enableSSH: true

delphix network interface 'vmxnet3s0' update addresses.0 *> set address=10.1.2.
4/24
delphix network interface 'vmxnet3s0' update addresses.0 *> get
  type: InterfaceAddress (*)
  address: 10.1.2.4/24 (*)
```

```
addressType: STATIC (*)
enableSSH: true (*)
```

5. Commit the operation.

```
delphix network interface 'vmxnet3s0' update addresses.0 *> commit
delphix network interface 'vmxnet3s0'> get
  type: NetworkInterface
  name: vmxnet3s0
  addresses:
    0:
      type: InterfaceAddress
      address: 10.1.2.4/24
      addressType: STATIC
      enableSSH: true
      state: OK
  dataNode: DATA_NODE-34
  device: vmxnet3s0
  macAddress: 0:c:29:32:96:a3
  mtu: 1500
  mtuRange: 60-9000
  reference: NETWORK_INTERFACE-vmxnet3s0-DATA_NODE-34
  state: OK
```

6. Re-enable the VDBs and dSources running from the engine.

CLI Cookbook: About Alert Notifications

The Delphix Engine can send out email notifications when alerts happen. Alert profiles control this functionality.

An alert profile is composed of two things:

- **Filter Specification:** A filter, or combination of filters, that specifies which alerts are of interest.
- **Alert Action:** This specifies the email addresses to which the Delphix Engine will send an email when an alert matches the filter specification.

By default, the Delphix Engine has a single alert profile configured with the following parameters:

- Filter Specification: Match any alert with a severity level of **CRITICAL** or **WARNING**.
- Alert Actions: Send an email to the address defined for user **admin**.

 The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username '`delphix_admin`'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Using the CLI, it is possible to:

- Modify the system default alert profile
- Create additional profiles in addition to the default one
- Set multiple actions for a single profile, such as "email `delphix_admin`" and "email `user1@mycompany.com`".

Simple Filters

- Filtered by Owner of alerts target – for example, objects owned by user 1

Complex Filters

Complex filters combine/modify other sub-filters:

- “And” filter – Used when all conditions defined must be met for the filter to notify the user with an email
- “Or” filter – Used when either one or the other of the conditions defined in the filters must be met for the filter to notify the user with an email
- “Not” filter – Used to exclude items

Limitations

- This is a CLI feature.
- Alert Profiles do not override permission settings. If you do not have Read permission on an object then your alert profile will never get triggered for that object's alerts, regardless of your filter settings.

The following CLI examples will run through how to create these three filters. Each example provides three different methods of setting up a profile. These include the following:

- A simple profile
- A profile with two filters
- A complicated profile

For more information, see [CLI Cookbook: Creating Alert Profiles](#).

A Simple Profile

A simple profile approach matches the Delphix out-of-the-box default alert profiles. To create a simple alert profile using the CLI as seen in the figure below, go into the alert profile section of the command-line interface (CLI) and create a new profile. Line four prompts the engine to send an email when the filters are triggered. The following three command lines refer to the filter specifications. Follow two severity levels: warning and critical. This will trigger an email alert when any warning or critical events occur.

```
twalsh-trunk.dcenter> cd alert
twalsh-trunk.dcenter alert> cd profile
twalsh-trunk.dcenter alert profile> create
twalsh-trunk.dcenter alert profile create *> set actions.0.type=AlertActionEmailUser
twalsh-trunk.dcenter alert profile create *> set filterSpec.type=SeverityFilter
twalsh-trunk.dcenter alert profile create *> set filterSpec.severityLevels.0=CRITICAL
twalsh-trunk.dcenter alert profile create *> set filterSpec.severityLevels.1=WARNING
twalsh-trunk.dcenter alert profile create *> commit
```

Simple Alert Profile example in the CLI

A Compound Alert Profile

Creating a compound alert profile in the CLI will combine two filters together. This profile triggers an email about any alert on objects owned by the delphix_admin plus any other alert that is critical. The compound alert profile creates two filters. The first one will be the target owner filter, which in this case is **admin**. The second filter is the severity filter, allowing users to match anything that is critical. Combine these two filters using the OR logic so that if any of the sub-filters match, the whole filter matches. An example of this can be seen in the figure below.



Alert Profile using OR logic

While working in the CLI, the first four lines describe a simple alert profile. The distinction between simple and compound alert profiles is that in a compound profile, the top-level filter uses an OR filter with sub-filters for target owner and severity level, as seen in line five of the figure below.


A Compound Alert Profile

Complex Alert Profile

A complex alert profile uses the profile filter created in the compound alert profile and modifies it. For the example shown in the figure below, you have a VDB named test_instance that you do not need any emails about. The following commands will create an effective filter.

1. Create an OR filter with two sub filters: target owner and event type.

2. Create a NOT filter which will exclude the VDB (test_instance) from which you do not want to receive notifications.
3. Use the AND logic to combine all these filters together, as seen below.



Complex Alert Profile

Below is an example of the command lines used to set up this complex profile.

```
twalsh-trunk.dcenter> cd alert
twalsh-trunk.dcenter alert> cd profile
twalsh-trunk.dcenter alert profile> create
twalsh-trunk.dcenter alert profile create *> set actions.0.type=AlertActionEmailUser
twalsh-trunk.dcenter alert profile create *> set filterSpec.type=AndFilter
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.0.type=NotFilter
twalsh-trunk.dcenter alert profile create *> edit filterSpec.subFilters.0.subFilter
twalsh-trunk.dcenter alert profile create filterSpec.subFilters.0.subFilter *> set type=TargetFilter
twalsh-trunk.dcenter alert profile create filterSpec.subFilters.0.subFilter *> set targets.0=test_instance
twalsh-trunk.dcenter alert profile create filterSpec.subFilters.0.subFilter *> back
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.1.type=OrFilter
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.1.subFilters.0.type=TargetOwnerFilter
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.1.subFilters.0.owners.0=delphix_admin
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.1.subFilters.1.type=SeverityFilter
twalsh-trunk.dcenter alert profile create *> set filterSpec.subFilters.1.subFilters.1.severityLevels=CRITICAL
twalsh-trunk.dcenter alert profile create *> commit
```

Complex Alert Profile CLI

CLI Cookbook: Creating Alert Profiles

This article describes how to create alert profiles.

Delphix generates alerts for different events. Users may want to be notified of events based on certain criteria such as the type of event or severity. An alert profile allows a user or group of users to be notified of the desired event.

⚠ The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Procedure

1. ssh into your engine using your `delphix_admin` username and password.

```
ssh admin@yourdelphixengine
```

2. Go into your alerts and list the alerts you already have.

```
delphix > alert
delphix alert > ls
```

3. Create your profile.

```
delphix alert > profile
delphix alert profile > create
delphix alert profile create * > ls
```

4. Set Actions and Severity Filter.

```
delphix alert profile create * > set actions.0.type=<AlertActionEmailList or
AlertActionEmailUser>
delphix alert profile create * > set actions.0.addresses.0=<email address to
send to>
delphix alert profile create * > set actions.0.addresses.1=<additional email
address>
delphix alert profile create * > set actions.0.addresses.2=<additional email
address>
delphix alert profile create * > ls
```

```
**Prior to Delphix Engine version 5.1
delphix alert profile create * > set severityFilter=
delphix alert profile create * > set severityFilter=<AUDIT|WARNING|CRITICAL|
INFORMATIONAL>
```

```
**Delphix Engine version 5.1 and beyond
set filterSpec.severityLevels=<AUDIT|WARNING|CRITICAL|INFORMATIONAL>
```

⚠ Valid Values for Parameters

actions.0.type:

- AlertActionEmailList: This type of alert is used to create an alert for any number of users. When this type is selected, an email address may be defined in each element of the "actions.0.addresses" array as illustrated above.
- AlertActionEmailUser: This type of alert is created for the email address of the user currently logged into the command-line interface. The "actions.0.addresses" array is not available for this type.

actions.0.filterSpec.severityLevels:

- INFORMATIONAL
- WARNING
- CRITICAL
- AUDIT

5. Commit your changes.

```
delphix alert profile create *> commit
```

Example:

```
ssh admin@yourengine
delphix > alert
delphix alert> ls
Objects
REFERENCE  TIMESTAMP                               TARGETNAME
EVENTTITLE
ALERT-102  2015-01-14T21:00:04.380Z  ASE/pubs2
Job complete
ALERT-101  2015-01-14T20:55:57.880Z  ASE/pubs2VDB
Job complete
ALERT-100  2015-01-14T19:35:32.958Z  ASE/pubs2VDB
Job complete
ALERT-99   2015-01-14T19:35:32.850Z  ASE/pubs2VDB
Job complete
ALERT-98   2015-01-14T19:34:58.744Z  ASE/pubs2
Error during job execution
ALERT-97   2015-01-14T18:12:01.928Z  ASE/pubs2
Job complete
ALERT-96   2015-01-14T18:03:10.664Z  ASE/pubs2
Job complete
ALERT-95   2015-01-14T17:16:07.464Z  ASE/pubs2
Job complete
ALERT-94   2015-01-14T17:15:55.298Z  ASE/market
Job complete
ALERT-93   2015-01-14T17:15:45.995Z  ASE/pubs2VDB
Job complete
ALERT-92   2015-01-14T16:39:33.133Z  nstacksolase2.acme.com-2015-01-14T16:39:13.821Z
Job complete
ALERT-91   2015-01-14T16:38:33.719Z  nstacksolase2.acme.com
Job complete
```

```

ALERT-90  2015-01-14T15:47:35.005Z  market
Validated sync failed for dSource
ALERT-89  2015-01-14T15:45:40.895Z  pubs2
Validated sync failed for dSource
ALERT-88  2015-01-14T15:02:14.874Z  ASE/market
Job complete
ALERT-87  2015-01-14T11:33:28.766Z  ASE/pubs2VDB
Job complete
ALERT-86  2015-01-13T23:11:46.838Z  ASE/market
Job complete
ALERT-85  2015-01-13T11:30:01.154Z  ASE/pubs2VDB
Job complete
ALERT-84  2015-01-13T11:07:04.385Z  pubs2
Backup detection failed
ALERT-83  2015-01-12T22:35:18.774Z  pubs2
Backup detection failed
ALERT-82  2015-01-12T11:30:00.063Z  ASE/pubs2VDB
Unable to connect to remote database during virtual database policy enforcement
ALERT-81  2015-01-12T11:30:00.054Z  ASE/pubs2
Unable to connect to remote database during dSource policy enforcement
ALERT-80  2015-01-12T08:38:26.983Z  pubs2
Backup detection failed
ALERT-79  2015-01-12T06:04:34.666Z  pubs2
Validated sync failed for dSource
ALERT-78  2015-01-11T11:30:03.393Z  ASE/pubs2VDB
Job complete
Children
profile
delphix alert> select ALERT-98
delphix alert "ALERT-98"> ls
Properties
  type: Alert
  event: alert.jobs.failed.object
  eventAction: Create the database on the target host.
  eventDescription: DB_EXPORT job for "ASE/pubs2" failed due to an error during
execution: Could not find database "pubs2VDB" on target instance "SRC_157_4K",
environment "ASE".
  eventSeverity: CRITICAL
  eventTitle: Error during job execution
  reference: ALERT-98
  target: ASE/pubs2
  targetName: ASE/pubs2
  targetObjectType: ASEDDBContainer
  timestamp: 2015-01-14T19:34:58.744Z
delphix alert> profile
delphix alert profile> select ALERT_PROFILE-1
delphix alert profile "ALERT_PROFILE-1"> ls
Properties
  type: AlertProfile
  actions:
    0:
      type: AlertActionEmailList
      addresses: sys_admin@acme.com

```

```

        format: HTML
    eventFilter: (empty)
    reference: ALERT_PROFILE-1
    severityFilter: CRITICAL,WARNING
    targetFilter: (empty)
Operations
delete
update
delphix alert profile> create
delphix alert profile create *> set actions.0.type=AlertActionEmailList
delphix alert profile create *> set actions.0.addresses.0=johndoe@acme.com
delphix alert profile create *> set actions.0.addresses.1=samsmith@acme.com
delphix alert profile create *> set filterSpec.severityLevels=INFORMATIONAL

```

*The last piece of the alert profile that needs to be configured is the "targetFilter". This is an array so you can define multiple targets. In the following example, there is a dSource named "pubs2" the user wants to define an alert on. If they try to set the filter to just the name of the dSource itself ("pubs2"), it will warn them that this is ambiguous and gives a hint on how to fully qualify it:

```

delphix > alert profile create
delphix alert profile create *> ls
Properties
  type: AlertProfile
  actions:
    0:
      type: AlertActionEmailList (*)
      addresses: foo@bar.com (*)
      format: HTML (*)
      filterSpec: (unset)
delphix alert profile create *> edit actions
delphix alert profile create actions *> add
delphix alert profile create actions 0 *> set addresses=foo@bar.com
delphix alert profile create actions 0 *> back
delphix alert profile create actions *> back
delphix alert profile create *> set filterSpec.severityLevels=WARNING
delphix alert profile create *> set filterSpec.type=TargetFilter
delphix alert profile create *> set filterSpec.targets="REPLICATION_SPEC_EXECUTE"
delphix alert profile create *> commit

```

Use the tab button freely to autocomplete and also see available options, for instance, while changing the severityLevels property, you can use the tab key like so:

```

DELPHIX-4221.dcenter alert profile 'ALERT_PROFILE-1' update *> set
filterSpec.severityLevels= <I HIT TAB HERE TO SEE OPTIONS BELOW>

```

```

AUDIT          CRITICAL          INFORMATIONAL  WARNING

```

 **Note on names used in the example**

SRC_157_4K: Repository (entity containing the database instances)

ASE: Group name

pubs2: Name of an individual database instance

The user sets the target filter to be equal to "pubs2/pubs2" and "ASE/pubs2" because if you review the "TARGETNAME" column from step 1 above, you will see alerts generated for both of these targets.

CLI Cookbook: Obtaining CPU Performance Information Using CLI

There are times when it may be desirable to obtain analytics information about the Delphix Engine CPU, which is formatted differently than what is available from the Delphix User Interface (UI).

Troubleshooting

You can obtain running information by looking at the **GUI > Resources > Performance Analytics**. You will see the graph for CPU if the CPU is selected.

When logged into the Delphix Engine command-line interface (CLI) as `delphix_admin`, you can see the default analytics gathered:

```

dlpx5120.dcenter> analytics
dlpx5120.dcenter analytics> list
NAME                STATISTICTYPE           STATE    COLLECTIONINTERVAL
COLLECTIONAXES

default.cpu         CPU_UTIL                RUNNING  1sec                idle,kernel,user
default.disk        DISK_OPS                RUNNING  1sec
op,avgLatency,latency,count,throughput
default.iscsi       iSCSI_OPS              RUNNING  1sec
op,latency,count,throughput
default.network     NETWORK_INTERFACE_UTIL  RUNNING  1sec
outBytes,networkInterface,inPackets,inBytes,outPackets
default.nfs         NFS_OPS                 RUNNING  1sec
op,latency,count,throughput
default.tcp         TCP_STATS               RUNNING  1sec
congestionWindowSize,localPort,remotePort,receiveWindowSize,inUnorderedBytes,sendWindow
Size,retransmittedBytes,outBytes,localAddress,roundTripTime,inBytes,unacknowledgedByt
es,remoteAddress

```

Resolution

You can also obtain some information from the CLI, using the [Performance Analytics Tool API Reference](#).

In order to view this information, you must log in with a user that had Delphix Admin privileges.

Option 1: Running an Existing Analytic

1. Login as admin and run `analytics`.

```

dlpx5120.dcenter> analytics
dlpx5120.dcenter analytics> list
NAME                STATISTICTYPE      STATE    COLLECTIONINTERVAL
COLLECTIONAXES

default.cpu         CPU_UTIL            RUNNING  1sec                idle,kernel,user
default.disk        DISK_OPS            RUNNING  1sec
op,avgLatency,latency,count,throughput
default.iscsi       iSCSI_OPS           RUNNING  1sec
op,latency,count,throughput
default.network     NETWORK_INTERFACE_UTIL  RUNNING  1sec
outBytes,networkInterface,inPackets,inBytes,outPackets
default.nfs         NFS_OPS             RUNNING  1sec
op,latency,count,throughput
default.tcp         TCP_STATS           RUNNING  1sec
congestionWindowSize,localPort,remotePort,receiveWindowSize,inUnorderedBytes,sendWindow
Size,retransmittedBytes,outBytes,localAddress,roundTripTime,inBytes,unacknowledgedByt
es,remoteAddress

Children
statistic

Operations
create

```

2. To see the available properties for CPU analytics, `select` and list the `default.cpu`.

```

dlpx5120.dcenter analytics> select default.cpu
dlpx5120.dcenter analytics 'default.cpu'> ls
Properties
  type: StatisticSlice
  name: default.cpu
  axisConstraints: (empty)
  collectionAxes: idle,kernel,user
  collectionInterval: 1sec
  dataNode: DATA_NODE-1
  reference: ANALYTICS_STATISTIC_SLICE-1
  state: RUNNING
  statisticType: CPU_UTIL

Operations
delete
getData
pause
rememberRange
resume
stopRememberingRange

```

3. Specify `setopt trace=true` in order to see the `CpuUtilDatapointStream` datapoints.

```
d\px5120.dcenter analytics 'default.cpu'> setopt trace=true

=== GET /resources/json/delphix/analytics/ANALYTICS_STATISTIC_SLICE-1 ===
=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": {
    "type": "StatisticSlice",
    "reference": "ANALYTICS_STATISTIC_SLICE-1",
    "namespace": null,
    "name": "default.cpu",
    "statisticType": "CPU_UTIL",
    "collectionInterval": 1,
    "state": "RUNNING",
    "collectionAxes": [
      "idle",
      "kernel",
      "user"
    ],
    "axisConstraints": [],
    "dataNode": "DATA_NODE-1"
  },
  "job": null,
  "action": null
}
=== END ===
```

4. Enter `getData` and `commit` to obtain the data gathered.

```
d\px5120.dcenter analytics 'default.cpu'> getData
=== GET /resources/json/delphix/analytics/ANALYTICS_STATISTIC_SLICE-1 ===
=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": {
    "type": "StatisticSlice",
    "reference": "ANALYTICS_STATISTIC_SLICE-1",
    "namespace": null,
    "name": "default.cpu",
    "statisticType": "CPU_UTIL",
    "collectionInterval": 1,
    "state": "RUNNING",
    "collectionAxes": [
      "idle",
      "kernel",
      "user"
    ],
  },
```

```

    "axisConstraints": [],
    "dataNode": "DATA_NODE-1"
  },
  "job": null,
  "action": null
}
=== END ===

```

```

dlpx5120.dcenter analytics 'default.cpu' getData *> commit
=== GET /resources/json/delphix/analytics/ANALYTICS_STATISTIC_SLICE-1/getData ===
=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": {
    "type": "DatapointSet",
    "resolution": 1,
    "datapointStreams": [
      {
        "type": "CpuUtilDatapointStream",
        "datapoints": [
          {
            "type": "CpuUtilDatapoint",
            "timestamp": "2016-12-06T13:53:30.000Z",
            "idle": 1946,
            "kernel": 33,
            "user": 19,
            "dtrace": null
          },
          {
            "type": "CpuUtilDatapoint",
            "timestamp": "2016-12-06T13:53:31.000Z",
            "idle": 1966,
            "kernel": 17,
            "user": 14,
            "dtrace": null
          },
          {
            "type": "CpuUtilDatapoint",
            "timestamp": "2016-12-06T13:53:32.000Z",
            "idle": 1968,
            "kernel": 17,
            "user": 13,
            "dtrace": null
          },
          {
            "type": "CpuUtilDatapoint",
            "timestamp": "2016-12-06T13:53:33.000Z",
            "idle": 1963,
            "kernel": 19,
            "user": 17,
            "dtrace": null
          }
        ]
      }
    ]
  }
}

```

```

        {
          "type": "CpuUtilDatapoint",
          "timestamp": "2016-12-06T13:53:34.000Z",
          "idle": 1968,
          "kernel": 16,
          "user": 15,
          "dtrace": null
        },
        ...
        ...
        {
          "type": "CpuUtilDatapoint",
          "timestamp": "2016-12-06T19:54:16.000Z",
          "idle": 1922,
          "kernel": 36,
          "user": 41,
          "dtrace": null
        },
        {
          "type": "CpuUtilDatapoint",
          "timestamp": "2016-12-06T19:54:17.000Z",
          "idle": 1953,
          "kernel": 26,
          "user": 20,
          "dtrace": null
        }
      ],
      "cpu": null
    }
  ],
  "overflow": false
},
"job": null,
"action": null
}
=== END ===
type: DatapointSet
datapointStreams:
  0:
    type: CpuUtilDatapointStream
    datapoints: [ ... ]
overflow: false
resolution: 1sec

```

Option 2: Creating a New Analytic

1. Log in as admin and then run `analytics`
2. Run `create` and `list` the properties.

```
dlpx5120.dcenter analytics> create
dlpx5120.dcenter analytics create *> ls
Properties
  type: StatisticSlice
  name: (required)
  axisConstraints: (unset)
  collectionAxes: (required)
  collectionInterval: (unset)
  dataNode: (unset)
  statisticType: (required)
```

3. Set the required Properties `name`, `collectionAxes` and `statisticType` (if you are unsure you can run `help <command>` for options)

```
set name=test.cpu
set collectionAxes=kernel
set statisticType=CPU_UTIL

dlpx5120.dcenter analytics create *> ls
Properties
  type: StatisticSlice
  name: test.cpu
  axisConstraints: (unset)
  collectionAxes: kernel
  collectionInterval: (unset)
  dataNode: (unset)
  statisticType: CPU_UTIL
```

4. `Commit` to begin data collection.

```
dlpx5120.dcenter analytics create *> commit

=== POST /resources/json/delphix/analytics ===
{
  "type": "StatisticSlice",
  "name": "test.cpu",
  "collectionAxes": [
    "kernel"
  ],
  "statisticType": "CPU_UTIL"
}
=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": "ANALYTICS_STATISTIC_SLICE-7",
  "job": null,
  "action": "ACTION-656"
}
=== END ===
`ANALYTICS_STATISTIC_SLICE-7
```

5. Use `ls` to see the new `test.cpu` analytic

```
dlpx5120.dcenter analytics> ls
```

NAME	STATISTICTYPE	STATE	COLLECTIONINTERVAL	
COLLECTIONAXES				
default.cpu	CPU_UTIL	RUNNING	1sec	idle, kernel, user
test.cpu	CPU_UTIL	RUNNING	1sec	kernel
default.disk	DISK_OPS	RUNNING	1sec	
op, avgLatency, latency, count, throughput				
default.iscsi	iSCSI_OPS	RUNNING	1sec	
op, latency, count, throughput				
default.network	NETWORK_INTERFACE_UTIL	RUNNING	1sec	
outBytes, networkInterface, inPackets, inBytes, outPackets				
default.nfs	NFS_OPS	RUNNING	1sec	
op, latency, count, throughput				
default.tcp	TCP_STATS	RUNNING	1sec	
congestionWindowSize, localPort, remotePort, receiveWindowSize, inUnorderedBytes, sendWindow Size, retransmittedBytes, outBytes, localAddress, roundTripTime, inBytes, unacknowledgedByt es, remoteAddress				
Children				
statistic				
Operations				
create				

6. Select `test.cpu` and use `getData` and `commit` to see the data gathered.

```
dlpx5120.dcenter analytics> select test.cpu
dlpx5120.dcenter analytics 'test.cpu'> ls
dlpx5120.dcenter analytics 'test.cpu'> getData
dlpx5120.dcenter analytics 'test.cpu' getData *> commit
```

7. Use `delete` to stop the collection and delete `test.cpu`

```

dlpx5120.dcenter analytics 'test.cpu'> delete
dlpx5120.dcenter analytics 'test.cpu' delete *> commit
=== POST /resources/json/delphix/analytics/ANALYTICS_STATISTIC_SLICE-7/delete ===
{}
=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": "",
  "job": null,
  "action": "ACTION-657"
}
=== END ===

dlpx5120.dcenter analytics> ls

NAME                STATISTICTYPE          STATE    COLLECTIONINTERVAL
COLLECTIONAXES

default.cpu          CPU_UTIL                RUNNING  1sec                idle,kernel,user
default.disk         DISK_OPS                RUNNING  1sec
op,avgLatency,latency,count,throughput
default.iscsi        iSCSI_OPS              RUNNING  1sec
op,latency,count,throughput
default.network      NETWORK_INTERFACE_UTIL  RUNNING  1sec
outBytes,networkInterface,inPackets,inBytes,outPackets
default.nfs          NFS_OPS                 RUNNING  1sec
op,latency,count,throughput
default.tcp          TCP_STATS               RUNNING  1sec
congestionWindowSize,localPort,remotePort,receiveWindowSize,inUnorderedBytes,sendWindow
wSize,retransmittedBytes,outBytes,localAddress,roundTripTime,inBytes,unacknowledgedByt
es,remoteAddress

Children
statistic

Operations
create
dlpx5120.dcenter analytics>

```

CLI Cookbook: Rebooting the Delphix Engine via CLI

Occasionally, the management stack will hang on the GUI, or routine maintenance will require a reboot of the Delphix Engine. You can do this safely through either the Delphix Setup or CLI. Before performing a reboot, it is important that all your VDBs are shut down and dSources disabled in order to maintain data integrity.

Prerequisites

- Shut down all VDBS
- Shut down dSources

Procedure

Complete the following steps to reboot the Delphix Engine via CLI:

1. Login to the CLI using the sysadmin **username** and **password**.

```
ssh sysadmin@yourdelphixengine
```

2. Go to **system > reboot**.

```
delphix > system  
delphix system > reboot
```

3. Commit the action.

```
delphix system reboot *> commit
```

CLI Cookbook: Disabling User-click Analytics

The Delphix User-click Analytics feature is a lightweight method to capture how users interact with Delphix product user interfaces. The goal of capturing this data is to get a better understanding of product usage, engagement, and user behavior, and to use this data to improve Delphix products and customer experience. This feature is enabled by default for customers deploying on or upgrading to this version. User-click Analytics may also be disabled via the UI.

 This procedure will disable user-click analytics on both the Delphix Engine and Delphix Self-Service.

Procedure to disable user-click analytics

1. ssh into your engine with a user that has Admin privileges.

```
ssh sysadmin@delphix
```

2. Go to Services and select the userInterface.

```
delphix > cd service  
delphix service > cd userInterface
```

3. Update the analyticsEnabled option to false.

```
delphix service userInterface > ls  
Properties  
  Type userInterfaceConfig  
  analyticsenabled: true  
Operations  
delphix service userInterface > update  
delphix service userInterface update *> set analyticsEnabled=false  
delphix service userInterface update *> commit
```

CLI Cookbook: Changing the API Version

This topic describes how to change the API version.

Procedure

1. Log in to the Delphix Engine as the admin user.

```
delphix> su - admin@DOMAIN
```

2. Run the version command.

```
delphix> version
1.10.3 # This lists the current API version.
delphix> version <version you want to set to>
delphix> version
1.8.2 # Displays the new version that you had set above.
```

CLI Cookbook: Hosts and Environments

These topics describe command-line interface procedures for working with database hosts and environments.

This section covers the following topics:

- [CLI Cookbook: Adding a UNIX Host](#)
- [CLI Cookbook: Adding a SQL Server Source Environment](#)
- [CLI Cookbook: Setting Multiple Addresses for a Target Host](#)
- [CLI Cookbook: How to Change Environment User](#)
- [CLI Cookbook: How to Create or Edit a Privilege Elevation Profiles and Profile Scripts](#)
- [CLI Cookbooks: Enabling and Configuring Environment Permissions](#)

CLI Cookbook: Adding a UNIX Host

This topic describes the process of adding a UNIX host using the 3.0 command-line interface.

Within Delphix, there are both hosts and host environments. A host represents a remote system, but may or may not be a source or target for linking or provisioning. For example, in an Oracle RAC cluster, the cluster environment represents the location of the Oracle installation(s), and while there are hosts within that cluster they are not individually manageable as environments.

Procedure

1. Create a new environment and set the parameter type to be a UNIX host.

```
delphix> environment create
delphix environment create *> set type=HostEnvironmentCreateParameters
delphix environment create *> set hostEnvironment.type=UnixHostEnvironment
delphix environment create *> set hostParameters.type=UnixHostCreateParameters
delphix environment create *> set
primaryUser.credential.type>PasswordCredential
delphix environment create *> get
  type: HostEnvironmentCreateParameters
  hostEnvironment:
    type: UnixHostEnvironment
    name: (unset)
    aseHostEnvironmentParameters: (unset)
    description: (unset)
    logCollectionEnabled: false
  hostParameters:
    type: UnixHostCreateParameters
    host:
      type: UnixHost
      address: (required)
      dspKeystoreAlias: (unset)
      dspKeystorePassword: (unset)
      dspKeystorePath: (unset)
      dspTruststorePassword: (unset)
      dspTruststorePath: (unset)
      javaHome: (unset)
      nfsAddressList: (unset)
      oracleJdbcKeystorePassword: (unset)
      privilegeElevationProfile: (unset)
      sshPort: 22
      sshVerificationStrategy: (unset)
      toolkitPath: (required)
  logCollectionEnabled: false
  primaryUser:
    type: EnvironmentUser
    name: (unset)
    credential:
      type: PasswordCredential
      password: (required)
```

```
environment: (unset)
groupId: (unset)
userId: (unset)
```

2. Set the host address.

The name of the environment is derived from the address used, though you can provide a more descriptive name if desired. The address can be DNS names, IP addresses, or a comma-separated list of the above.

```
delphix environment create *> set hostParameters.host.address=192.168.1.2
```

3. Set the toolkit path.

This is where Delphix will store temporary binaries used while the host is configured as part of Delphix.

```
delphix environment create *> set hostParameters.host.toolkitPath=/work
```

4. Set the username and password to use when connecting over SSH.

This user must have the privileges described in the Delphix Administration Guide. To configure an SSH user, change the credential type to `SystemKeyCredential`.

```
delphix environment create *> set primaryUser.name=oracle
delphix environment create *> set primaryUser.credential.password
Enter primaryUser.credential.password: *****)
```

5. Commit the result.

The environment discovery process will execute as an asynchronous job. The default behavior is to wait for the result, so progress will be updated until the discovery process is complete or fails.

```
delphix environment create *> commit
UNIX_HOST_ENVIRONMENT-4
Dispatched job JOB-39
ENVIRONMENT_CREATE_AND_DISCOVER job started for "192.168.1.2".
ENVIRONMENT_CREATE_AND_DISCOVER job for "192.168.1.2" completed
successfully.
delphix>
```

CLI Cookbook: Adding a SQL Server Source Environment

This topic describes how to add a SQL Server source environment using the command line interface.

Since SQL Server source environments do not have the Delphix Connector running on them, you must use a target environment as a proxy when adding source environments. Delphix uses the connector running on the proxy environment to run commands against the source environment. See [Overview of Requirements for SQL Server Environments](#) for more information.

Procedure

Enter these commands through the command-line interface:

```
/environment;

create;
set type=HostEnvironmentCreateParameters;

set hostEnvironment.type=WindowsHostEnvironment;
set hostEnvironment.name=<Source environment name>;
set hostEnvironment.proxy=<target host name>;

set hostParameters.type=WindowsHostCreateParameters;
set hostParameters.host.type=WindowsHost;
set hostParameters.host.address="<Source host IP address or hostname>";

set primaryUser.name="<domain\username>";
set primaryUser.credential.type>PasswordCredential;
set primaryUser.credential.password=<password>;

commit;
```

Example

The CLI commands for adding source host "mssql_source_1" using target host "mssql_target_1" as proxy and environment user "ad\delphix_user" would be:

```
/environment;

create;
set type=HostEnvironmentCreateParameters;

set hostEnvironment.type=WindowsHostEnvironment;
set hostEnvironment.name="mssql_source_1";
set hostEnvironment.proxy="mssql_target_1";

set hostParameters.type=WindowsHostCreateParameters;
set hostParameters.host.type=WindowsHost;
set hostParameters.host.address="mssql_source_1";

set primaryUser.name="ad\delphix_user";
```

```
set primaryUser.credential.type=PasswordCredential;  
set primaryUser.credential.password="i_am_the_password";  
  
commit;
```

CLI Cookbook: Setting Multiple Addresses for a Target Host

This topic is an example of using arrays to configure a target host to support multiple IP addresses. The `nfsAddressList` property is an array of strings.

Procedure

1. Select the host to update

```
delphix> host
delphix host> select example
delphix host "example"> update
```

2. Set the address:

```
Delphix host '192.168.121.141' update *> set nfsAddressList="192.168.1.23,192.168.2.44"
```

3. Get the current addresses, both as a string and as an array object.

```
delphix host "example" update *> get nfsAddressList
  192.168.1.23,192.168.2.44 (*)
delphix host "example" update *> get nfsAddressList[0]
  192.168.1.23 (*)
delphix host "example" update *> edit nfsAddressList
delphix host "example" update addresses *> get
  0:
    192.168.1.23 (*)
  1:
    192.168.2.44 (*)
```

4. Commit the result:

```
delphix host "example" update addresses *> commit
delphix host "example">
```

CLI Cookbook: How to Change Environment User

1. ssh into your engine using Admin privileges

```
ssh admin@delphix
```

2. Go to Environment and find the Environment you would like to update

```
delphix > environment
delphix environment > ls
delphix environment > select test_env
```

3. Select Environment updating and Update

```
delphix environment "test_env" > update
delphix environment "test_env" update *> ls
```

4. Set `primaryUser` to new user you would like to use for the Environment (NOTE: new user should be added as an environment user before making as a primaryUser)

```
delphix environment "test_env update" *> set primaryUser=<new user>
```

5. Commit the change

```
delphix environment "test_env" update *> commit
```

Example:

```
ssh admin@delphix
delphix > environment
delphix environment > ls

Objects

NAME      DESCRIPTION

Demo

Children

oracle

user
```

Operations

create

```
delphix environment > select Demo
delphix environment "Demo" > update
delphix environment "Demo" update *> ls
```

Properties

type: UnixHostEnvironment

name: Demo

description:

primaryUser: delphix

```
delphix environment "Demo" update *> set primaryUser=<new user>
delphix environment "Demo" update *> commit
```

CLI Cookbook: How to Create or Edit a Privilege Elevation Profiles and Profile Scripts

Background:

If you are running a version prior to 6.0.0 please contact your Professional Services representative.

Procedure for Creating an Elevation Profile:

1. Log into the CLI using `delphix_admin` or a user with Admin privileges and got to 'Host'.

```
ssh admin@youengine
youengine > host
```

2. Select `privilegeElevation` then `profile`.

```
youengine host > privilegeElevation
youengine host privilegeElevation > profile
```

3. Set the name of the profile and the version of the profile.

```
youengine host privilegeElevation profile *> set name=<profilename>
youengine host privilegeElevation profile *> set version=<profileversion>
```

4. Commit the profile to save it.

```
youengine host privilegeElevation profile *> commit
```

Procedure for creating a profileScript

Please note that you will need to create the script yourself or with the help of the Professional Services team.

1. Log into the CLI using `delphix_admin` or a user with Admin privileges and got to 'Host'.

```
ssh delphix_admin@youengine
youengine > host
```

2. Select `privilegeElevation` then `profileScript`.

```
youengine host > privilegeElevation
youengine host privilegeElevation > profileScript
```

3. Create your script by setting, name, contents, and profile (this can be your previously created profile or the default sudo).

```
youengine host privilegeElevation profileScript > create
youengine host privilegeElevation profileScript *> set name=<scriptname>
youengine host privilegeElevation profileScript *> set contents=<your script>
```

```
youengine host privilegeElevation profileScript *> set profile=<yourprofile>
```

4. Commit to save.

```
youengine host privilegeElevation profileScript *> commit
```

Privilege Elevation Profiles and Delphix Replication

Because all Delphix Engines have a default Profile called `sudo`, which would normally exist on both source and target Delphix Engines, replication collisions which would normally prevent a successful failover are automatically resolved.

- Only Profiles which are actually assigned to a host are replicated. All currently unassigned profiles are ignored.
- Profile name collisions are resolved by the display names of duplicate Profiles being prefixed with a unique object identifier. This is described in more detail in the next section.
- A default Profile will not retain default status after replication failover. The Profile assigned as default on the replication target Delphix Engine will remain the default. Therefore, if the source Delphix Engine has a non-standard default Profile, it will need to be manually set as the new default on the replication target Delphix Engine after failover.

Caveats

By design, the Delphix Engine allows the creation of Profiles with duplicate names. This is not a bug. It exists for several reasons:

- This allows replication failover to complete without duplicate Profile names triggering a collision.
- Makes versioning possible so that a profile with the same name can have multiple versions as iterations are made (some of which may not be production-ready).

However, this behavior has the consequence of changing the display name of Profiles. Once a duplicate name exists, a unique object identifier is prefixed to the name. Any references to such a profile (such as assigning to a host) must use the long format with the unique identifier. For example, standard Delphix Engine has the following Profile:

```
Delphix5031HWv8> host privilegeElevation profile ls
Objects
NAME      ISDEFAULT
sudo      true
Operations
create
```

If a new profile that is also called "`sudo`" is created, the display names automatically change as follows:

```
Objects
NAME                                     ISDEFAULT
`HOST_PRIVILEGE_ELEVATION_PROFILE-1/sudo  true
`HOST_PRIVILEGE_ELEVATION_PROFILE-5/sudo  false
Operations
create
```

Known issues

- It is not possible to delete a Profile. Attempting to do so results in an API error. However, Profiles can be renamed to something meaningful like "`unused_1`".

- Profiles created that contain single quotation marks can no longer be selected. They become orphaned Profiles.
- When pasting in script contents, the cursor does not correctly move to the end of the last line being pasted. Unless the cursor is moved to the end of the line before pressing ENTER, the script will not be complete.

Neither of the above issues has any operational impact.

CLI Cookbooks: Enabling and Configuring Environment Permissions

This topic describes how to enable the environment permissions feature to restrict what users can do with environments.

By default, all engine users can list all environments and hosts and see their details. Moreover, all users are able to link dSources from and provision VDBs to any environment without requiring any permissions on environments, as long as they have [appropriate permissions on the target group where the dsource or VDB will be located](#).

Enabling Environment and Permissions

To restrict non-administrator users from seeing, linking from, and provisioning to any environment, Delphix administrators can enable environment authorizations.

```
delphix> authorization configuration
delphix authorization configuration > ls
Properties
  type: AuthorizationConfig
  environmentAndHostAuth: false

Operations
update
delphix authorization configuration> update
delphix authorization configuration update *> set environmentAndHostAuth=true
delphix authorization configuration update *> commit
```

Similarly, to go back to the default state in which all users have permission to perform those operations, the Delphix administrator must set the `environmentAndHostAuth` property back to `false`.

Granting and Revoking Permissions on Environments and Hosts

When environment permissions are enabled, only Delphix administrators can list environments and hosts, see their details, or link dSources from or provision VDBs to environments.

To authorize any other user to perform such an operation on an environment or host, a Delphix administrator must create an appropriate authorization.

```
delphix> authorization create
delphix authorization create *> set user=someuser
delphix authorization create *> set role=PROVISIONER
delphix authorization create *> set target=SourceEnvironment:/somehost.example.com
```

To revoke an authorization, a Delphix administrator must delete the corresponding authorization object.

```
delphix> authorization
delphix> ls
REFERENCE          USER          ROLE          TARGET
```

```

AUTHORIZATION-1 sysadmin OWNER sysadmin
AUTHORIZATION-2 admin OWNER admin
AUTHORIZATION-3 admin OWNER domain0
AUTHORIZATION-4 someuser Data SourceEnvironment:/somehost.example.com

delphix authorization> select `AUTHORIZATION-4
delphix authorization '(USER-2, ROLE-2, UNIX_HOST_ENVIRONMENT-1)'> delete
delphix authorization '(USER-2, ROLE-2, UNIX_HOST_ENVIRONMENT-1)' delete *> commit

```

Permissions on Environments and Hosts

Role	Environment Privileges	Host Privileges
Owner	<ul style="list-style-type: none"> • Can provision VDBs from the environment • Can link dSources from the environment • Can access the same information as a Reader 	<ul style="list-style-type: none"> • Can access the same information as a Reader
Provisioner	<ul style="list-style-type: none"> • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption • Can provision VDBs from owned dSources and VDBs 	<ul style="list-style-type: none"> • Can access the same information as a Reader
Data Operator	<ul style="list-style-type: none"> • Can access statistics on the dSource, VDB, or snapshot such as usage, history, and space consumption • Can refresh or rollback VDBs • Can snapshot dSources and VDBs 	<ul style="list-style-type: none"> • Can access the same information as a Reader
Reader	<ul style="list-style-type: none"> • Can see the configuration of the environment 	<ul style="list-style-type: none"> • Can see the configuration of the host
Self-Service Only	<ul style="list-style-type: none"> • Can access the same information as a Reader 	<ul style="list-style-type: none"> • Can see the configuration of the host

CLI Cookbook: Source Databases and dSources

These topics describe command-line interface procedures for working with dSources.

This section covers the following topics:

- [CLI Cookbook: Detaching and Attaching a SQL Server dSource](#)
- [CLI Cookbook: Disabling LogSync for a dSource](#)
- [CLI Cookbook: Enabling Oracle Validated Sync](#)
- [CLI Cookbook: Linking a SQL Server Database Loading from a Specific Full Backup of the Source Database](#)
- [CLI Cookbook: Linking a SQL Server Database Loading from the Last Full Backup of the Source Database](#)
- [CLI Cookbook: Linking to a Single Instance Oracle Database](#)
- [CLI Cookbook: Listing Data Source Sizes](#)
- [CLI Cookbook: Detaching and Attaching an Oracle dSource](#)
- [CLI Cookbook: How to Change Database User Password](#)
- [CLI Cookbook: Changing an SAP ASE dSource's Staging Database](#)
- [CLI Cookbook: Detaching and Attaching a SAP ASE dSource](#)
- [CLI Cookbook: Linking an SAP ASE Database Loading from the Last Full Backup of the Source Database](#)

CLI Cookbook: Detaching and Attaching a SQL Server dSource

This topic describes how to detach and attach a SQL server dSource. When attaching a SQL Server dSource to a new data source, the new data source must be in the same database satisfying the following constraints:

- pptRepository needs to be set to the name of the SQL instance on the staging server. The unlink operation removes the database from the SQL instance on the staging server and unmounts the iscsi luns, reattaching the dSource via the CLI will remount the iscsi luns and puts the database back.

Prerequisites

A dSource can only be attached to a new data source once it has been [unlinked](#).

Procedure

1. Select dSource.

```
delphix> database "dexample"
```

2. Run the detachSourcecommand, specifying the current active source. This step can be skipped if the dSource has already been detached. through the GUI.

```
delphix database "dexample"> detachSource
delphix database "dexample" detachSource *> set source=dexample
delphix database "dexample" detachSource *> commit
```

3. Run the `attachSource` command.

```
delphix database "dexample"> attachSource
```

4. Set the following for SQL Server:

You can also type `help pptRepository` to see what is wanted

You can also set `pptRepository=<`then press `tab`> to list all values.

```
delphix database "dexample" attachSource *> edit attachData
delphix database "dexample" attachSource attachData *> set type=MSSqlAttachData
delphix database "dexample" attachSource attachData *> set config=SQLSERVER/dexample
delphix database "dexample" attachSource attachData *> set sharedBackupLocations="<\
\SERVER1\Backups>"
delphix database "dexample" attachSource attachData *> set pptRepository=SQL2008R2
delphix database "dexample" attachSource attachData *> set
mssqlUser.type=MSSqlDomainUser
delphix database "dexample" attachSource attachData *> set mssqlUser.user=ad\dbuser
delphix database "dexample" attachSource attachData *> set
mssqlUser.password=password=dbuserpwd
delphix database "dexample" attachSource attachData *> edit ingestionStrategy
```

```
delphix database "dexample" attachSource attachData ingestionStrategy *>set
type=ExternalBackupIngestionStrategy
delphix database "dexample" attachSource attachData ingestionStrategy *>set
validatedSyncMode=TRANSACTION_LOG
delphix database "dexample" attachSource attachData ingestionStrategy *> back
delphix database "dexample" attachSource attachData *> edit operations.preSync
delphix database "dexample" attachSource attachData operations.preSync *> back
delphix database "dexample" attachSource attachData*> edit operations.postSync
delphix database "dexample" attachSource attachData operations.postSync *> back
```

5. Commit the operation.

```
delphix database "dexample" attachSource *> commit
```

CLI Cookbook: Disabling LogSync for a dSource

This topic provides a simple example of how the nested state is represented and manipulated. The LogSync state is maintained in the sourcingPolicy property of dSources, itself an object with several different fields.

Procedure

1. Select the dSource to be changed and run the `update` command.

```
delphix> database "example"  
delphix "example"> update
```

2. Get the current property using dot-delimited notation.

```
delphix "example" update *> get sourcingPolicy.logsyncEnabled  
true
```

3. The property could also be set using dot-delimited notation, but for illustrative purposes, we can also use the `edit` command and set it directly.

```
delphix "example" update *> edit sourcingPolicy  
delphix "example" update sourcingPolicy *> set logsyncEnabled=false
```

4. Commit the state, either from within the editing context or after running `back` to return to the parent context.

```
delphix "example" update sourcingPolicy *> commit  
delphix "example">
```

CLI Cookbook: Enabling Oracle Validated Sync

This topic describes how to designate a staging host and enable validated sync for Oracle data sources.

Prerequisite - Designating a Staging Host

In order to validate an Oracle dSource snapshot for syncing, the Delphix Engine requires a host with an Oracle installation that is compatible with the dSource. This machine is known as the **staging** host. You must explicitly designate which machines you want the Delphix Engine to use as staging hosts. All machines that have been marked as staging sites are added to a pool. During validated sync, the Delphix Engine will select a compatible host from the pool, export the requisite archived redo logs and datafiles, and execute Oracle media recovery on the host. Follow these steps to designate a staging host.

1. Select the repository you want to designate as staging.

```
delphix>/repository/select '/u01/app/ora10205/product/10.2.0/db_1'
```

2. Execute the `update` command.

```
delphix repository "/u01/app/ora10205/product/10.2.0/db_1">update
```

3. Set staging to `true`.

```
delphix repository "/u01/app/ora10205/product/10.2.0/db_1" update *>set staging=true
```

4. `Commit` the operation to designate the repository as staging.

```
delphix repository "/u01/app/ora10205/product/10.2.0/db_1" update *> commit
```

To configure validated sync for multiple dSources with different Oracle versions, you must designate a compatible staging source for each. If multiple compatible staging sites exist, the Delphix Engine will select one at random.

 The validated sync process will consume some resources on the staging host when snapshots are taken. Designating a performance-critical host as a staging host is not recommended.

Procedure - Enabling Validated Sync

1. Select the dSource for which you want to enable validated sync.

```
delphix>/database/select redsox1
```

2. Execute the `update` command.

```
delphix database "redsox1">update
```

3. Set `preProvisioningEnabled` to `true`.

```
delphix database "redsox1" update *>set preProvisioningEnabled=true
```

4. `Commit` the operation to enable validated sync.

```
delphix database "redsox1" update *>commit
```

CLI Cookbook: Linking a SQL Server Database Loading from a Specific Full Backup of the Source Database

This topic describes how to use the command-line interface to link a SQL Server database by loading from a specific full backup of the source database as indicated by the backup UUID.

Prerequisites

- You can get the backup UUID for the all backup files of a chosen database using the following query on the source database under the column **backup_set_uuid**

```
Use master
select backupset.database_name,
       backupset.type,
       backupset.backup_set_id,
       backupset.backup_set_uuid,
       backupset.family_guid,
       backupset.position,
       backupset.first_lsn,
       backupset.last_lsn,
       backupset.database_backup_lsn,
       backupset.name,
       backupset.has_bulk_logged_data,
       backupset.is_damaged,
       backupset.begins_log_chain,
       backupset.is_copy_only,
       backupset.backup_finish_date,
       backupset.database_version,
       backupset.database_guid,
       mediafamily.logical_device_name,mediafamily.physical_device_name
from msdb.dbo.backupmediafamily mediafamily join msdb.dbo.backupset backupset
on mediafamily.media_set_id = backupset.media_set_id where backupset.database_name =
N'<Database Name>'
order by backupset.backup_finish_date desc
```

Procedure

Enter these commands through the Delphix Engine command-line interface:

```
/database;

link;
set type=LinkParameters;
set name=<dSource name>;
set group=<group name>;

set linkData.type=MSSQLLinkData;
set linkData.syncParameters.type=MSSQLExistingSpecificBackupSyncParameters;
set linkData.config=<source database>;
set linkData.sharedBackupLocations="<source database backup locations>";
set linkData.pptRepository=<SQL instance on the staging server>;
set linkData.sourcingPolicy.type=SourcingPolicy;
set linkData.mssqlUser.type=MSSQLDomainUser
```

```
set linkData.mssqlUser.user=ad\dbuser
set linkData.mssqlUser.password.password=dbuserpwd
set linkData.syncParameters.backupUUID=<backup UUID>;
set linkData.ingestionStrategy.type=<ingestion strategy type>;
set linkData.ingestionStrategy.validatedSyncMode=<validated sync mode type>;

commit;
```

CLI Cookbook: Linking a SQL Server Database Loading from the Last Full Backup of the Source Database

This topic describes how to use the command-line interface to link a SQL Server database by loading from the last full backup of the source database.

Procedure

Enter the following commands in the Delphix Engine command-line interface:

```
/database;  
  
link;  
set type=LinkParameters;  
set name=<dSource name>;  
set group=<group name>;  
  
set linkData.type=MSSqlLinkData;  
set linkData.syncParameters.type=MSSqlExistingMostRecentBackupSyncParameters  
set linkData.config=<source database>;  
set linkData.sharedBackupLocations="<source database backup locations>";  
set linkData.pptRepository=<SQL instance on the staging server>;  
set linkData.sourcingPolicy.type=SourcingPolicy;  
set linkData.mssqlUser.type=MSSqlDomainUser  
set linkData.mssqlUser.user=ad\dbuser  
set linkData.mssqlUser.password.password=dbuserpwd  
set linkData.ingestionStrategy.type=<ingestion strategy type>;  
  
commit;
```

CLI Cookbook: Linking to a Single Instance Oracle Database

This topic describes how to link to a single instance Oracle database using the Delphix Engine command-line interface.

Prerequisites

You will need the following information:

- The name of the dSource you want to create.
- The group in which you want to create the dSource.
- The database unique name of the Oracle database you want to link to.
- The database username/password with sufficient privileges as described in the Delphix User Guide.
- The host environment user with sufficient privileges as described in the Delphix User Guide.

Procedure

1. Execute the `database link` command.

```
delphix> database link
delphix database link>
```

2. The default `linkData.type` in the `LinkParameters` is set to `ASELinkData`, but you can confirm that by getting the input type. Set `linkData.type` to `OracleLinkFromExternal` for Oracle database.

```
delphix database link *> get linkData.type
  ASELinkData
delphix database link *> set linkData.type=OracleLinkFromExternal
delphix database link *> get linkData.type
  OracleLinkFromExternal (*)
delphix database link *> ls
Properties
  type: LinkParameters
  name: (required)
  description: (unset)
  group: (required)
  linkData:
    type: OracleLinkFromExternal (*)
    backupLevelEnabled: (unset)
    bandwidthLimit: (unset)
    checkLogical: (unset)
    compressedLinkingEnabled: (unset)
    config: (required)
    dbCredentials:
      type: PasswordCredential
      password: (required)
    dbUser: (required)
    diagnoseNoLoggingFaults: (unset)
    encryptedLinkingEnabled: (unset)
    environmentUser: (required)
```

```

externalFilePath: (unset)
filesPerSet: (unset)
linkNow: (unset)
nonSysCredentials: (unset)
nonSysUser: (unset)
numberOfConnections: (unset)
operations: (unset)
preProvisioningEnabled: (unset)
rmanChannels: (unset)
sourcingPolicy: (unset)
syncParameters:
  type: OracleSyncFromExternalParameters
  doNotResume: (unset)
  doubleSync: (unset)
  forceFullBackup: (unset)
  skipSpaceCheck: (unset)

```

Operations
defaults

3. Set the name for the dSource and the group in which you want to create it.

```

delphix database link *> set name=example1
delphix database link *> set group="<New Group>"

```

4. Set the source configuration.

For Oracle databases, these are identified by the database unique name. If you are unsure of the set of available databases, you can list available source configurations.

```

delphix database link *> /sourceconfig list
NAME      REPOSITORY      LINKINGENABLED
example1  '/opt/ora/dexample1' true
example2  '/opt/ora/dexample1' true
delphix database link *> set linkData.config=example1

```

5. Set the privileged database username/password.

The password can be set like other properties, or the value can be omitted so that it can be manually inputted without exposing the password.

```

delphix database link *> set linkData.dbUser=delphix
delphix database link *> set linkData.dbCredentials.password
Enter dbCredentials.password: *****

```

6. Set the privileged environment user.

This user must be from the same environment as the associated source config set in step 4. You can list the set of available users through the `environment user list` command.

```

delphix database link *> /environment/user list
NAME

```

```
oracle
delphix database link *> set linkData.environmentUser=oracle
```

- Adjust any other properties you may want, such as RMAN tunables, description, and whether to link now. The full set of options is described in the API documentation for the `OracleLinkFromExternal` type. If you set the `linkNow` property, then this operation will wait for the sync to complete, otherwise, you can perform the initial link by running the sync command at a later point.

```
delphix database link *> set linkData.linkNow=true
```

- Commit the result.

```
delphix database link *> commit
`ORACLE_DB_CONTAINER-1
Dispatched job JOB-8
DB_LINK job started for "<New Group>/example1".
Obtaining information from source database "<New Group>/example1".
Creating new TimeFlow for dSource "<New Group>/example1".
The dSource "example1" was successfully linked from source database "<New
Group>/example1".
DB_LINK job for "<New Group>/example1" completed successfully.
delphix>
```

CLI Cookbook: Listing Data Source Sizes

This topic describes a basic use of the CLI `list` command.

1. Switch to the source view and view the default list.

```
delphix> source
delphix source> list
NAME      CONTAINER  VIRTUAL  CONFIG
example   example   false    example
vexample  vexample  true     vexample
```

2. List sources with their database size.

```
delphix> source
delphix source> select example
delphix source 'example'> get runtime.databaseSize
  2.80GB
```

CLI Cookbook: Detaching and Attaching an Oracle dSource

This topic describes how to attach a dSource to a different data source.

Prerequisites

Before detaching an Oracle dSource, you must capture the following RMAN configuration:

- Level or SCN based backups
- Data load channel settings: number of channels and files per channel

The above RMAN configuration will be removed once the dSource is unlinked.

A dSource can only be attached to a new data source once it has been unlinked.

When attaching an Oracle dSource to a new data source, the new data source must be the same logical database satisfying the following constraints:

- Same dbid
- Same dbname
- Same creation time
- Same resetlogs SCN
- Same resetlogs time
- Same redo stream, where a log must exist with
 - Same sequence
 - Same thread
 - Same end SCN

For Oracle dSources, this procedure can be used to initially link from a standby server that is faster or less disruptive, unlink the dSource, and then attach it to the production server for subsequent incremental SnapSync operations. When you perform the attach operation, you will need the source config name of an unlinked database.

Procedure

1. Select **dSource**.

```
delphix> database "dexample"
```

2. Run the **detachSource** command, specifying the currently active source. This step can be skipped if the dSource has already been detached through the GUI.

```
delphix database "dexample"> detachSource
delphix database "dexample" detachSource *> set source=name-of-old-src-DB-
server
delphix database "dexample" detachSource *> commit
```

3. Run the **attachSource** command.

```
delphix database "dexample"> attachSource
```

4. Set the config to point to an unlinked source.

The following is an example to attach to an Oracle data source: (Hint: use <TAB> to complete variable names and known values)

```

delphix database "dexample" attachSource *> set
attachData.type=OracleAttachData
delphix database "dexample" attachSource *> set attachData.config=name-of-
dSource-as-shown-in-environment
delphix database "dexample" attachSource *> set
attachData.environmentUser=myuser
delphix database "dexample" attachSource *> set
attachData.oracleFallbackUser=orauser
delphix database "dexample" attachSource *> edit
attachData.oracleFallbackCredentials
delphix database "dexample" attachSource *> set set type>PasswordCredential
delphix database "dexample" attachSource *> set password=orauserpwd

```

The attachData.config listing can be observed in the administration GUI, under Manage (pull-down at the top) --> Environments --> choose your dSource --> select Databases (upper right).

The dSources shown on the new source DB server may acquire odd names eg. "P02:UNKNOWN:vwxyz". These are a side-effect of having multiple instances of the same container name (possible in DR); Delphix needs to disambiguate. If these do show, they indicate containers in the new environment. Successfully attaching to these and deleting the old environment will re-establish the original names.

5. Commit the operation.

```

delphix database "dexample" attachSource *> commit

```

CLI Cookbook: How to Change Database User Password

1. ssh into your engine using Admin privileges.

```
ssh admin@delphixengine
```

2. Go to sourceconfig and find the Database that you need to update the password on.

```
delphix > sourceconfig
delphix sourceconfig > ls
delphix sourceconfig > select <yourdatabase>
```

3. Update the password.

```
delphix sourceconfig "yourdatabase" > update
delphix sourceconfig "yourdatabase" update *> ls
delphix sourceconfig "yourdatabase" update *> set credentials.password=<new password>
```

4. Commit the change.

```
delphix sourceconfig "database" update *> commit
```

Example:

```
ssh admin@example
delphix > sourceconfig
delphix sourceconfig > ls

Objects

NAME      REPOSITORY                LINKINGENABLED
meta1    '/u01/oracle/10.2.0.4/ee1' true

Operations
create

delphix sourceconfig > select meta1
delphix sourceconfig "meta1" > ls

Properties
  type: OracleSIConfig
  name: meta1
  credentials:
```

```
    type: PasswordCredential
    password: ****
databaseName: meta1
discovered: true
environmentUser: delphix
instance:
  type: OracleInstance
  instanceName: meta1
  instanceNumber: 1
linkingEnabled: true
nonSysCredentials: (unset)
nonSysUser: (unset)
reference: ORACLE_SINGLE_CONFIG-1
repository: '/u01/oracle/10.2.0.4/ee1'
services:
  0:
    type: OracleService
    discovered: true
    jdbcConnectionString: jdbc:oracle:thin:@172.16.100.69:1525:meta1
  1:
    type: OracleService
    discovered: true
    jdbcConnectionString: jdbc:oracle:thin:@172.16.100.69:1521:meta1
uniqueName: meta1
user: delphix
```

Operations

delete

update

validateCredentials

```
delphix sourceconfig "meta1" > update
```

```
delphix sourceconfig "meta1" update *> set credentials.password=<new password>
```

```
delphix sourceconfig "meta1" update *> commit
```

CLI Cookbook: Changing an SAP ASE dSource's Staging Database

Prerequisites

In order to change an SAP ASE dSource's staging database, you need to know the name of the staging database and you need to disable the dSource.

Procedure

1. To find the name of the staging database, hover the mouse over the "Staging Database" name on the dSource card in the GUI or issue the following command via the CLI (replacing "pubs2" with the name of the dSource and "delphix.acme.com" with the hostname of your Delphix Engine):

```
$ echo "/source; select pubs2; ls" | ssh admin@delphix.acme.com | grep stagingSource
Password:
stagingSource: dxpb91a7LJlwj933xvLPvd_pubs2
```

2. Disable the dSource by either toggling the "Enable/Disable" toggle in the Delphix Management application or by using the CLI:

```
$ ssh admin@delphix.acme.com
delphix> source
delphix source> select pubs2
delphix source 'pubs2'> disable
delphix source 'pubs2' disable *> set type=SourceDisableParameters
delphix source 'pubs2' disable *> commit
    Dispatched job JOB-365
    SOURCE_DISABLE job started for "pubs2".
    SOURCE_DISABLE job for "pubs2" completed successfully.
```

3. Change the repository of the staging database to reside in the desired SAP ASE instance:

```
$ ssh admin@delphix.acme.com
delphix> sourceconfig
delphix sourceconfig> select dxpb91a7LJlwj933xvLPvd_pubs2
delphix sourceconfig 'dxpb91a7LJlwj933xvLPvd_pubs2'> update
delphix sourceconfig 'dxpb91a7LJlwj933xvLPvd_pubs2' update *> ls
Properties
  type: ASESICConfig
  credentials: (unset)
  databaseName: dxpb91a7LJlwj933xvLPvd_pubs2
  environmentUser: nstacksolasetest/sybase
  linkingEnabled: false
  repository: nstacksolasetest/SRC_157_4K
  user: (unset)
delphix sourceconfig 'dxpb91a7LJlwj933xvLPvd_pubs2' update *> set
  environmentUser=nstacksolasetest/sybase
delphix sourceconfig 'dxpb91a7LJlwj933xvLPvd_pubs2' update *> set repository=
```

```

nstacksolaseprod/SRC_157_4K                                NSTACK_16K
RH68_ASE157_S1
SQL2008R2                                                  SQL2012
nstacksolasestg
nstacksolasetest/SRC_157_4K
delphix sourceconfig 'dxbp91a7LJlwj933xv1Pvd_pubs2' update *> set
  repository=nstacksolasestg
delphix sourceconfig 'dxbp91a7LJlwj933xv1Pvd_pubs2' update *> commit

```

Hint: Rapidly hit the "tab" key twice after typing "set repository=" to make the CLI present a list of available instances.

4. Enable the dSource:

```

delphix sourceconfig 'dxbp91a7LJlwj933xv1Pvd_pubs2'> /source
delphix source> select pubs2
delphix source 'pubs2'> enable
delphix source 'pubs2' enable *> set type=SourceEnableParameters
delphix source 'pubs2' enable *> commit
  Dispatched job JOB-367
  SOURCE_ENABLE job started for "pubs2".
  SOURCE_ENABLE job for "pubs2" completed successfully.

```

CLI Cookbook: Detaching and Attaching a SAP ASE dSource

This CLI cookbook recipe describes how to Detach and Attach an SAP ASE dSource using the CLI.

Prerequisites

A dSource can only be attached to a new data source once it has been unlinked.

When attaching an SAP ASE dSource to a new data source, the new data source must be the same logical database satisfying the following constraints:

- Same dbid
- Same dbname
- Same creation time

You must also make sure that you follow the normal prerequisites for an SAP ASE data source found in [SAP ASE Support and Requirements](#).

Procedure

Detach a dSource

1. Login to the CLI as admin or a user with Admin privileges.
2. Select dSource.

```
delphix> database "dexample"
```

3. Run the `detachSource` command, specifying the currently active source. Note: This step can be skipped if the dSource has already been detached through the GUI.

```
delphix database "dexample"> detachSource
delphix database "dexample" detachSource *> set source=dexample
delphix database "dexample" detachSource *> commit
```

Attach a dSource

1. Login to the CLI as admin or a user with Admin privileges.
2. Run the `attachSource` command.

```
delphix database "dexample"> attachSource
delphix database "dexample" attachSource *> set attachData.config=dexample
delphix database "dexample" attachSource *> set
attachData.dbCredentials.password=sybase
delphix database "dexample" attachSource *> set attachData.dbUser=sa
delphix database "dexample" attachSource *> set attachData.loadBackupPath=/tmp/
backups
delphix database "dexample" attachSource *> set attachData.sourceHostUser="source_hos
t_environment/sybase"
delphix database "dexample" attachSource *> set attachData.stagingHostUser="staging_h
ost_environment/sybase"
```

```
delphix database "dexample" attachSource *> set attachData.stagingRepository="staging
_ASE_servername_example"
```

```
delphix database "dexample" attachSource *> ls
Properties
```

```
type: ASEAttachSourceParameters
```

```
attachData:
```

```
type: ASEAttachData
```

```
config: dexample (*)
```

```
dbCredentials:
```

```
type: PasswordCredential
```

```
password: ***** (*)
```

```
dbUser: sa (*)
```

```
dumpCredentials: (unset)
```

```
externalFilePath: (unset)
```

```
loadBackupPath: /tmp/backups (*)
```

```
loadLocation: (unset)
```

```
mountBase: (unset)
```

```
operations: (unset)
```

```
sourceHostUser: source_ASE_servername_example/sybase (*)
```

```
stagingHostUser: staging_ASE_servername_example/sybase (*)
```

```
stagingPostScript: (unset)
```

```
stagingPreScript: (unset)
```

```
stagingRepository: staging_ASE_servername_example (*)
```

```
validatedSyncMode: ENABLED
```

```
delphix database "dexample" attachSource *> commit
```

```
ASE_DB_CONTAINER-3
```

```
Dispatched job JOB-25
```

```
DB_ATTACH_SOURCE job started for "Untitled/dexample".
```

DB_ATTACH_SOURCE job **for** "Untitled/dexample" completed successfully.

 This command is only necessary if you are using a Remote Backup Server configuration from staging to the source, instead of an NFS mounted shared directory for backups and transaction log dumps:

```
delphix database "dexample" attachSource *> set  
attachData.loadLocation.backupServerName=source_backupserver_name_example
```

CLI Cookbook: Linking an SAP ASE Database Loading from the Last Full Backup of the Source Database

This topic describes how to use the command-line interface to link an SAP ASE database by loading from the most recent full backup of the source database.

Procedure

Enter the following commands in the Delphix Engine command-line interface:

```
ssh delphix_admin@delphix
Password:
delphix> /database link
delphix database link *> set linkData.type=ASELinkData
delphix database link *> set name=db2
delphix database link *> set group="ASE dSource"
delphix database link *> set linkData.config=db2
delphix database link *> set linkData.dbUser=sa
delphix database link *> set linkData.dbCredentials.password=sybase
delphix database link *> set linkData.loadBackupPath=/mnt/dump
delphix database link *> set linkData.mountBase=/mnt/provision/
vdb_or_staging_database_name
delphix database link *> set linkData.sourceHostUser=nstackrh69/sybase
delphix database link *> set linkData.stagingHostUser=nealorarh75/sybase
delphix database link *> set linkData.stagingRepository=ASE157SP138
delphix database link *> set linkData.sourcingPolicy.logsyncEnabled=true
delphix database link *> set linkData.sourcingPolicy.type=SourcingPolicy
delphix database link *> unset linkData.syncParameters
delphix database link *> edit linkData.syncParameters
delphix database link linkData.syncParameters *> back
delphix database link *> ls
Properties
  type: LinkParameters
  name: db2 (*)
  description: (unset)
  group: ASE dSource (*)
  linkData:
    type: ASELinkData
    config: db2 (*)
    dbCredentials:
      type: PasswordCredential
      password: ***** (*)
    dbUser: sa (*)
    dumpCredentials: (unset)
    dumpHistoryFileEnabled: false
    externalFilePath: (unset)
    loadBackupPath: /mnt/dump (*)
    loadLocation: (unset)
    mountBase: /mnt/provision/vdb_or_staging_database_name (*)
    operations: (unset)
    sourceHostUser: nstackrh69/sybase (*)
    sourcingPolicy:
      type: SourcingPolicy (*)
```

```

    logsyncEnabled: true (*)
    stagingHostUser: nealorarh75/sybase (*)
    stagingOperations: (unset)
    stagingPostScript: (unset)
    stagingPreScript: (unset)
    stagingRepository: ASE157SP138 (*)
    syncParameters:
      type: ASELatestBackupSyncParameters (*)
    validatedSyncMode: ENABLED

```

Operations

defaults

```
delphix database link *> commit
```

```
`ASE_DB_CONTAINER-209
```

```
Dispatched job JOB-3704
```

```
DB_LINK job started for "ASE dSource/db2".
```

```
DB_LINK job for "ASE dSource/db2" completed successfully.
```

mountBase option

The **mountBase** parameter is an option that was added in Delphix 5.2 and higher. By default, Delphix mounts the staging database under the Delphix toolkit directory. If a directory is specified for this parameter, the staging database's NFS devices will be mounted under this directory rather than under the toolkit directory. This can be especially helpful when the dSource is linked with LogSync enabled. SAP ASE has a limit of 127 characters for the fully qualified path to the transaction logs specified in the "LOAD TRANSACTION" statement. When LogSync is enabled, Delphix keeps a copy of the transaction logs on the engine under the dSources "archive" folder. The default naming convention for the folders under the toolkit can easily cause the 127 character limit to be exceeded so it is highly recommended to use this parameter when enabling LogSync. The mountBase is limited to 87 characters (the device names Delphix generates are 32 characters and the subdirectory containing the archive files is about 8 characters long). This leaves approximately 40 characters for the name of the transaction logs themselves. The mountBase parameter must be unique for each VDB or staging database. The path can reside under a common parent directory for example, /mnt/provision but you must specify a unique child directory under the parent for each VDB or staging database using this optional parameter.

CLI Cookbook: VDBs

These topics describe command-line interface procedures for working with virtual databases (VDBs).

This section covers the following topics:

- [CLI Cookbook: Attaching or Detaching a PDB](#)
- [CLI Cookbook: Attaching, Detaching, or Linking a CDB](#)
- [CLI Cookbook: Changing SGA Parameter](#)
- [CLI Cookbook: Changing the SID of Oracle RAC VDBs](#)
- [CLI Cookbook: Toggle new DBID generation upon refresh options for Oracle VDBs](#)
- [CLI Cookbook: Creating a Policy](#)
- [CLI Cookbook: Creating a VDB Config Template](#)
- [CLI Cookbook: Determining the Snapshot used to provision a VDB](#)
- [CLI Cookbook: How to Refresh a VDB from a specific Snapshot](#)
- [CLI Cookbook: Oracle VDB Migration](#)
- [CLI Cookbook: Provisioning a SAP ASE VDB](#)
- [CLI Cookbook: Provisioning a Single Instance Non-multitenant Oracle VDB](#)
- [CLI Cookbook: Provisioning a SQL Server VDB](#)
- [CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark](#)
- [CLI Cookbook: Provisioning a Virtual PDB](#)
- [CLI Cookbook: Provisioning a Virtual PDB in a Target CDB](#)
- [CLI Cookbook: Refresh a VDB from a Specific Timepoint or Latest](#)
- [CLI Cookbook: Repairing a Timeflow](#)
- [CLI Cookbook: Rolling Back a VDB](#)
- [CLI Cookbook: Rolling Forward a VDB](#)
- [CLI Cookbook: Taking a Snapshot](#)
- [CLI Cookbook: V2P: Virtual to Physical of a Single Instance Non-multitenant Oracle Database](#)
- [CLI Cookbook: V2P \(Virtual to Physical\) of a Single Instance Oracle Database with Datafiles on Separate File Systems](#)
- [CLI Cookbook: V2P Virtual to Physical on SQL Server](#)
- [CLI Cookbook: VDB status](#)
- [CLI Cookbook: Provisioning a Virtual PDB from a Non-Multitenant Source Database](#)
- [CLI Cookbook: Migrating a Virtual PDB and a Virtual CDB](#)
- [CLI Cookbook: Migrating an Oracle RAC Virtual PDB and a Virtual CDB](#)
- [CLI Cookbook: Locating and Updating the Value of tdeKeyIdentifier](#)

CLI Cookbook: Attaching or Detaching a PDB

PDB CLI Attach

1. Ensure the new source environment has been fully discovered by the Delphix Engine, including CDB discovery.
2. Login to the engine via SSH with a Delphix admin account.
3. Select the PDB dSource database to be attached.

```
dvde.dcenter> database
dvde.dcenter database> select "R268PDB1"
dvde.dcenter database 'R268PDB1'>
```

4. Attach the PDB dSource to the correct source PDB in the new environment using the attachSource command; you will need to specify the environment user and login credentials for the PDBdSource.

```
dvde.dcenter database 'R268PDB1'> attachSource
dvde.dcenter database 'R268PDB1' attachSource *> set
attachData.type=OraclePDBAttachData
dvde.dcenter database 'R268PDB1' attachSource *> set attachData.config=R268PDB1
dvde.dcenter database 'R268PDB1' attachSource *> set
attachData.environmentUser=ora12201
dvde.dcenter database 'R268PDB1' attachSource *> set attachData.dbUser=delphix
dvde.dcenter database 'R268PDB1' attachSource *> set
attachData.dbCredentials.password
Enter attachData.dbCredentials.password: *****
dvde.dcenter database 'R268PDB1' attachSource *> commit
R268PDB1
Dispatched job JOB-46
DB_ATTACH_SOURCE job started for "Untitled/R268PDB1".
Obtaining information from source database "Untitled/R268PDB1".
The dSource "R268PDB1" was successfully linked from source database
"Untitled/R268PDB1".
```

PDB CLI Detach

1. Login to the engine via SSH with a Delphix admin account.
2. Select the PDB dSource database to be detached.

```
dvde.dcenter> database
dvde.dcenter database> select "R268PDB1"
dvde.dcenter database 'R268PDB1'>
```

3. Detach the PDB dSource using the detachSource command:

```
dvde.dcenter database 'R268PDB1'> detachSource
dvde.dcenter database 'R268PDB1' detachSource *> set source=R268PDB1
dvde.dcenter database 'R268PDB1' detachSource *> commit
```

```
Dispatched job JOB-45  
DB_DETACH_SOURCE job started for "Untitled/R268PDB1".  
DB_DETACH_SOURCE job for "Untitled/R268PDB1" completed successfully.
```

CLI Cookbook: Attaching, Detaching, or Linking a CDB

This topic describes how to attach, detach, or link a CDB using the command-line interface.

A CDB is not automatically detached when you detach the last PDB of the CDB. The CDB still remains in an inactive state which in turn prevents you from removing the environment. Therefore, you must detach the CDB in order to remove the environment.

Note

You can perform detach or attach operations only via Command-Line Interface (CLI).

CDB CLI Detach

Perform the following steps to detach a CDB.

Note

You can detach a CDB only when there are no PDBs linked and the CDB does not have any vPDBs.

1. Login to the engine via SSH with a Delphix admin account.

```
$ ssh admin@YOUR_ENGINE
```

2. Select the CDB dSource database to be detached.

```
dvde.dcenter> cd database
dvde.dcenter database> select CDOMLOSR421F
dvde.dcenter database 'CDOMLOSR421F'>
```

3. Detach the CDB dSource using the detachSource command.

```
dvde.dcenter database 'CDOMLOSR421F'> detachSource
dvde.dcenter database 'CDOMLOSR421F' detachSource *> set source=CDOMLOSR421F
dvde.dcenter database 'CDOMLOSR421F' detachSource *> commit
  Dispatched job JOB-45
  DB_DETACH_SOURCE job started for "Untitled/CDOMLOSR421F".
  DB_DETACH_SOURCE job for "Untitled/CDOMLOSR421F" completed successfully.
```

You must attach the CDB back before performing the following operations: linking a new PDB, attaching a PDB, provisioning a vPDB, or attaching a converted PDB. Failure to do this will either result in an error due to duplicated object or will create a second CDB that will duplicate the space used for the CDB.

CDB CLI Link

You can not perform an attach operation for PDBs that have already been detached from the primary CDB and are required to be attached to a standby CDB if the standby CDB is not already linked or attached. You must link the CDB manually before performing the PDB attach operation.

Perform the following steps to link a CDB.

1. Login to the engine via SSH with a Delphix admin account.

```
$ ssh admin@YOUR_ENGINE
```

2. Select the CDB dSource database to be linked.

```
dvde.dcenter> cd database
dvde.dcenter database> link
dvde.dcenter database link *>
```

3. Link the CDB dSource using the link command.

```
dvde.dcenter database link *> set name=mycdb
dvde.dcenter database link *> set group=Untitled
dvde.dcenter database link *> edit linkData
dvde.dcenter database link linkData *> set type=OracleLinkFromExternal
dvde.dcenter database link linkData *> set syncStrategy.config=CDOMSHSR6706
dvde.dcenter database link linkData *> set environmentUser=oracle
dvde.dcenter database link linkData *> set linkNow=true
dvde.dcenter database link linkData *> commit
`ORACLE_DB_CONTAINER-23
Dispatched job JOB-171
DB_LINK job started for "Untitled/mycdb".
Obtaining information from source database "Untitled/mycdb".
Creating new TimeFlow for dSource "Untitled/mycdb".
Request to SnapSync container database "Untitled/mycdb" after attaching it
will be ignored.
Action: Direct SnapSync of container database is not allowed. Container
database SnapSync will be taken when SnapSync of corresponding
pluggable database is taken.
The dSource"mycdb" as successfully linked from source database "Untitled/
mycdb".
DB_LINK job for "Untitled/mycdb" completed successfully.
```

CDB CLI Attach

Perform the following steps to attach a CDB.

1. Ensure the new source environment has been fully discovered by the Delphix Engine.
2. Login to the engine via SSH with a Delphix admin account.

```
$ ssh admin@YOUR_ENGINE
```

3. Select the CDB dSource database to be attached.

```
dvde.dcenter> cd database
dvde.dcenter database> select CDOMLOSR421F
dvde.dcenter database 'CDOMLOSR421F'>
```

4. Attach the CDB dSource to the correct source CDB in the new environment using the attachSource command.

Note

When attaching a CDB linkNow=true command is ignored, then the CDBs can not be snapshotted directly. CDB snapshots will be taken when PDB requires a snapshot of the CDB.

```
dvde.dcenter database 'CDOMLOSR421F'> attachSource
dvde.dcenter database 'CDOMLOSR421F' attachSource *> edit attachData
dvde.dcenter database 'CDOMLOSR421F' attachSource attachData *> set
type=OracleAttachData
dvde.dcenter database 'CDOMLOSR421F' attachSource attachData *> set
config=CDOMLOSR421F
dvde.dcenter database 'CDOMLOSR421F' attachSource attachData *> set
environmentUser=oracle
dvde.dcenter database 'CDOMLOSR421F' attachSource *> commit
CDOMLOSR421F
Dispatched job JOB-46
DB_ATTACH_SOURCE job started for "Untitled/CDOMLOSR421F".
Starting validation of attach parameters for database "CDOMLOSR421F".
Obtaining information from source database "Untitled/CDOMLOSR421F".
The dSource "CDOMLOSR421F" was successfully linked from source database
"Untitled/CDOMLOSR421F".
```

CLI Cookbook: Changing SGA Parameter

Below outlines the procedure to change SGA parameter setting on a provisioned VDB.

Procedure

1. Log into the Delphix Management application as `admin` or a user with Admin privileges.
2. Go to `source` and then `select` the name of the VDB that you would like to change the parameters of.
3. You are then going to `update` and edit the `configParams`.
4. Next, you are going to set `sga_target=` correct value.
5. `Commit` the operation so that it saves.

Example

```
ssh admin@engine
delphix > source
delphix source > select "vdb_example"
delphix source "vdb_example" > update
delphix source "vdb_example" *> edit configParams
delphix source "vdb_example" *> set sga_target=new value
delphix source "vdb_example" *> commit
```

 Modifying configuration parameters for a vPDB is not supported, so this workflow will not succeed for a vPDB.

CLI Cookbook: Changing the SID of Oracle RAC VDBs

This topic describes how to change the SID of instances in an Oracle RAC VDB.

This example demonstrates how to switch the instance name and number between two different hosts, from

```
SQL> select * FROM V$ACTIVE_INSTANCES;
```

```
INST_NUMBER INST_NAME
```

```
-----
```

```
1 cnrac3:VchiBEB1
```

```
2 cnrac4:VchiBEB2
```

to

```
SQL> select * FROM V$ACTIVE_INSTANCES;
```

```
INST_NUMBER INST_NAME
```

```
-----
```

```
1 cnrac4:VchiBEB1
```

```
2 cnrac3:VchiBEB2
```

Procedure

1. Stop the VDB through the GUI and login to the Delphix CLI.
2. Select the sourceconfig of the RAC VDB whose instances you would like to rename.

```
kfc-manual.dcenter> sourceconfig
kfc-manual.dcenter sourceconfig> select Vchicago_BEB
```

3. Use the update command to change the properties of the sourceconfig.

```
kfc-manual.dcenter sourceconfig "Vchicago_BEB"> update
```

4. Type 'ls' to view the complete list of properties associated with the VDB's sourceconfig. For configurations with larger numbers of RAC instances, the listing may not show the individual instances but will instead display [...]. In order to see the instance configuration, type 'edit instances'.

```
kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> ls
```

```
Properties
```

```
  type: OracleRACConfig
```

```
  credentials:
```

```
    type: PasswordCredential
```

```
    password: *****
```

```
  environmentUser: ora1024
```

```
  instances:
```

```
    0:
```

```
      type: OracleRACInstance
```

```
      instanceName: VchiBEB1
```

```
      instanceNumber: 1
```

```

    node: cnrac4
  1:
    type: OracleRACInstance
    instanceName: VchiBEB2
    instanceNumber: 2
    node: cnrac3
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: '/u01/app/ora1024/product/10.2.0/db_1'
  services: [ ... ]
  user: delphix

```

5. Use the Set command to change the values for instanceName and instanceNumber for each instance.

```

kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> set instances.0.instanceName
=VchiBEB2
kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> set instances.0.instanceNumb
er=2
kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> set instances.1.instanceName
=VchiBEB1
kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> set instances.1.instanceNumb
er=1

```

6. Finally, commit the changes.

```

kfc-manual.dcenter sourceconfig "Vchicago_BEB" update *> commit;

```

7. Restart the VDB through the GUI for the changes to take effect on the VDB.

CLI Cookbook: Toggle new DBID generation upon refresh options for Oracle VDBs

This topic describes how to toggle **Generate new DBID upon refresh** option for Oracle VDBs.

Procedure

1. Login to the Delphix CLI.
2. Go to source

```
delphix> source
```

3. Select the VDB that you need to update

```
delphix source> select 'VDBOMSRE71EE4_QGE'
```

4. Toggle new DBID value.

The new DBID value can be true or false. If the `newDBID` parameter is set to false then the new DBID generation will stop.

```
delphix source VDBOMSRE71EE4_QGE> update  
delphix source VDBOMSRE71EE4_QGE> set newDBID=<new value>
```

5. Finally, commit the changes to save the changes.

```
delphix source VDBOMSRE71EE4_QGE> commit
```

6. Refresh the VDB through CLI to generate a new DBID if the `newDBID` parameter was set to true. If the `newDBID` parameter was set to false, the Delphix engine will use the same DBID as the parent.

```
delphix> /database  
delphix database> select 'VDBO_QGE'  
delphix database 'VDBO_QGE'> refresh  
delphix database 'VDBO_QGE' refresh> commit
```

CLI Cookbook: Creating a Policy

This will outline how to create a policy in the CLI, please note that you can also do this in the GUI.

Procedure

1. ssh into your Delphix Engine using admin credentials.

```
ssh admin@delphixengine
delphix > ls
```

2. Go to `policies` and `createAndApply` (please note that you cannot just create a policy, you must `createAndApply`, in the GUI you have the option to just create) and set your policy parameter.

```
delphix > policy
delphix policy > createAndApply
delphix policy createAndApply * > set policy.type=< choose from QuotaPolicy,
RefreshPolicy, RetentionPolicy, SnapshotPolicy or SyncPolicy)
delphix policy createAndApply * > set policy.name=< name your policy>
delphix policy createAndApply * > set policy.customized=true
delphix policy createAndApply * > set policy.
delphix policy createAndApply * > set policy.provisionSource=(LATEST_SNAPSHOT or
LATEST_TIME_FLOW_LOG)
```

-  *if doing a `RefreshPolicy`, `SyncPolicy` or `SnapshotPolicy` you are also going to need to add the following:

```
delphix policy createAndApply * > edit policy.scheduleList
delphix policy createAndApply policy.scheduleList * > add
delphix policy createAndApply policy.scheduleList * > set cronString=
delphix policy createAndApply policy.scheduleList * > set cutoffTime=
delphix policy createAndApply policy.scheduleList * > back
```

3. Set your target parameters which are going to be a container, group etc.

```
delphix policy createAndApply * > set target=
```

4. Verify and `commit`.

```
delphix policy createAndApply * > ls
delphix policy createAndApply * > commit
```

CLI Cookbook: Creating a VDB Config Template

This topic will address how to create a VDB Config Template in the CLI; this functionality is also available in the GUI.

Procedure

1. ssh into your Delphix Engine using admin credentials.

```
ssh admin@<yourdelphixengine>
```

2. Go to `database` and then `template` and then `create`.

```
delphix > database template
delphix database template > create
delphix database template create *> set name=<set the name of database
template>
delphix database template create *> set parameters.<set parameters you want>
delphix database template create *> set sourceType=<set the source>
```

3. Verify information and commit.

```
delphix database template create *> ls
delphix database template create *> commit
```

Example:

```
ssh admin@testengine
testengine > database template
testengine database template > create
testengine database template create *> set name=oracle
testengine database template create *> set parameters.PGA_TARGET=100M
testengine database template create *> set parameters.MEMORY_TARGET=100M
testengine database template create *> set sourceType=OracleVirtualSource
testengine database template create *> ls
```

Properties

```
type: DatabaseTemplate
name: oracle (*)
description: (unset)
parameters:
  MEMORY_TARGET: 100M (*)
  PGA_TARGET: 100M (*)
sourceType: OracleVirtualSource (*)
testengine database template create *> commit
```

CLI Cookbook: Determining the Snapshot used to provision a VDB

Procedure:

The parent snapshot can be determined using the CLI as follows:

1. Log into the server as a Delphix administrator:

```
ssh admin@<server_ip>
```

2. Select the VDB.

```
delphix> database
delphix database> ls
Objects
NAME          PARENTCONTAINER DESCRIPTION
dSource1      -
dSource2      -
VDB1          dSource1      -
VDB2          dSource2      -
VDB3          dSource1      -
delphix database> select VDB1
```

3. List the VDB parameters, and make a note of the currentTimeflow value.

```
delphix database "VDB1"> ls
Properties
  type: OracleDatabaseContainer
  name: VDB1
  currentTimeflow: VDB1/default
  description: (unset)
  diagnoseNoLoggingFaults: true
  endianness: BIG_ENDIAN
  group: <New Group>
  masked: false
  os: HP-UX
  parentContainer: dSource1
  performanceMode: false
  processor: ia64
  reference: ORACLE_DB_CONTAINER-10
  runtime:
    type: OracleDBContainerRuntime
    logSyncActive: true
  sourcingPolicy:
    type: OracleSourcingPolicy
    encryptedLinkingEnabled: false
    logsyncEnabled: true
    logsyncInterval: 300
    logsyncMode: ARCHIVE_ONLY_MODE
  version:
```

4. Select the Timeflow listed for the VDB.

```
delphix database "VDB1"> /timeflow
delphix timeflow> select VDB1/default
List the timeflow parameters. The Snapshot used to provision the VDB is listed
as parentSnapshot

delphix timeflow "VDB1/default"> ls
Properties
```

5. List the Timeflow parameters. The Snapshot used to provision the VDB is listed as parentSnapshot.

```
delphix timeflow "VDB1/default"> ls
Properties

type: OracleTimeflow
name: VDB1/default
container: VDB1
parentPoint:
  type: OracleTimeflowPoint
  location: 141285148
  timeflow: dSource1/default
parentSnapshot: @2013-02-14T15:07:28.491Z
reference: ORACLE_TIMEFLOW-92572
```

CLI Cookbook: How to Refresh a VDB from a specific Snapshot

These steps will allow you to refresh a VDB from any specific snapshot, not just the most recent one.

1. Identify the VDB and snapshot that you want to use.

```
ssh admin@<yourengine>
delphix > database ls
delphix database > /snapshot
delphix snapshot > list database=<SOURCEOFSNAPSHOT>
```

2. Go to the `database` and `refresh`.

```
delphix > /database
delphix database > refresh
```

3. Now set what type of refresh you are going to do.

```
delphix database 'VDB' refresh *> set
timeflowPointParameters.type=TimeflowPointSnapshot
```

4. Set the snapshot.

```
delphix database 'VDB' refresh *> set timeflowPointParameters.snapshot=@XXXX-
XX-XXTX:XX:XX.XXXZ
```

5. Commit the action.

```
delphix database 'VDB' refresh *> commit
```



You can use tab to complete most actions in the CLI in addition to listing the possibilities that are available when setting parameters.

CLI Cookbook: Oracle VDB Migration

This topic describes moving a VDB from one environment or installation to another.

VDBs can be moved (or migrated) between hosts by changing the source repository associated with the VDB source config.

Restrictions

The following restrictions apply when migrating a VDB between repositories:

- When migrating a RAC VDB, the host of each `OracleRACInstance` must be updated as well.
- The mount point of the VDB source cannot be changed.
- The `database_unique_name` and `db_name` cannot be changed.
- The new environment and repository must be a compatible target environment.

Procedure

1. Select the source associated with the VDB. By default, sources are named the same as the VDB.

```
delphix> source "vexample"
```

2. Disable the source by running the `disable` command and committing the operation.

```
delphix source "vexample"> disable
delphix source "vexample" disable *> commit
  Dispatched job JOB-171
  SOURCE_DISABLE job started for "vexample".
  Starting disable of virtual database.
  Unexporting storage.
  Virtual database disable successful.
  SOURCE_DISABLE job for "vexample" completed successfully.
delphix source "vexample">
```

3. Select the source config associated with the source. By default this is also the same name as the VDB.

```
delphix source "vexample"> get config
  vexample
delphix source "vexample"> /sourceconfig "vexample"
delphix sourceconfig "vexample">
```

4. Update the repository and repository user associated with the source config.

 You must use the Environment name of the Environment because the repository requires the environment name and oracle home location.

```
delphix sourceconfig "vexample"> update
delphix sourceconfig "vexample" update *> set repository=192.168.100.247/'/opt/oracle/product/10.2.0.4/db_1'
```

```
delphix sourceconfig "vexample" update *> set environmentUser=192.168.100.247/
ora1024
delphix sourceconfig "vexample" update *> commit
delphix sourceconfig "vexample">
```

5. Enable the source.

```
delphix sourceconfig "vexample"> /source "vexample" enable
delphix source "vexample" enable *> commit
  Dispatched job JOB-18
  SOURCE_ENABLE job started for "vexample".
  Enabling dataset "vexample".
  Exporting storage containers from the Delphix Engine.
  Mounting datasets.
  Mounting filesystems for the virtual database instance "1".
  Starting virtual database.
  Starting instance 1 on virtual database "vexample".
  Virtual database "vexample" was successfully started.
  Dataset "vexample" enabled.
  SOURCE_ENABLE job for "vexample" completed successfully.
delphix sourceconfig "vexample">
```

CLI Cookbook: Provisioning a SAP ASE VDB

This topic describes how to provision an SAP ASE VDB using the command line interface.

Prerequisites

You will need the following information:

- The name of the VDB you want to create
- The group in which to create the VDB
- The SAP ASE database name for the VDB
- The source dSource or VDB from which you wish to provision
- The semanticLocation, LSN, or timestamp of the point you want to provision from (if not using the most recent). You can run these commands to get the list of snapshots or timeflow ranges:

```
snapshot list database=dexample
snapshot list timeflow=dexample
snapshot list fromDate="2020-03-01T00:00:00.000Z" toDate="2020-03-04T11:31:27.883Z"
```

- The target host on which you want to create the VDB. You can list the hosts with the /host list command.
- The source repository (SAP ASE instance on the target host) in which to create the VDB. These can be listed with the /repository list command.

Procedure

1. Execute the database provision command.

```
delphix> database provision
```

2. Set the type for the new VDB

```
delphix database provision *> set type=ASEProvisionParameters
```

3. Use defaults to fill in most of the information and then customize any additional information that you do not want defaulted, for what information has been filled in after defaults you can do an ls for all fields:

```
delphix database provision *> defaults
delphix database provision *> ls
Properties:
  type: TimeflowPointSemantic
  container: (unset)
  location: LATEST_POINT
delphix database provision *> set container=<dexample>
delphix database provision *> commit
```

4. Set the name and group for the new VDB.

```
delphix database provision *> set container.name=<vexample>
```

```
delphix database provision *> set container.group="<New Group>"
```

5. Set the name of the new VDB.

```
delphix database provision *> set sourceConfig.databaseName=<vexample>
```

6. Set the target Dataset Home.

```
delphix database provision *> set sourceConfig.repository=<Dataset Home>  
delphix database provision *> set sourceConfig.environmentUser=<Host  
environment name/sybase>
```

7. Set the source container from which to provision.

```
delphix database provision *> set timeflowPointParameters.container=<dexample>
```

8. Set the desired value for truncateLogOnCheckpoint

```
delphix database provision *> set truncateLogOnCheckpoint=false
```

9. Commit the configuration and start the DB_PROVISION job

```
delphix database provision *> commit
```

CLI Cookbook: Provisioning a Single Instance Non-multitenant Oracle VDB

This topic describes how to provision a single instance non-multitenant Oracle VDB using the Delphix Engine command-line interface.

Prerequisites

You will need the following information:

- The name of the VDB you want to create
- The group in which to create the VDB
- The Oracle database name
- The Oracle database unique name
- The Oracle database instance number
- The Oracle database instance name
- The source dSource or VDB from which you wish to provision. This will be referenced as the "container" in the "defaults" command below.
- The semanticLocation, SCN, or timestamp of the point you want to provision from. You can run these commands to get the list of snapshots or Timeflow ranges:

```
snapshot list database=dexample
timeflow "dexample" timeflowRanges; commit
```

- The base mount point on the target server where VDB data should be mounted
- The source repository (oracle install) in which to create the VDB. These can be listed with the `repository list` command.
- If you are using a VDB template, the name of the template to use.

Procedure

1. Execute the `database provision` command.

```
delphix> database provision
```

2. Execute the `defaults` command. Once you commit this command, it will return a partially constructed provision parameters object.

```
delphix database provision> defaults
```

3. Set the Timeflow point source Timeflow and location.

```
delphix database provision defaults *> set type=TimeflowPointSemantic
delphix database provision defaults *> set container=dexample
delphix database provision defaults *> set location=LATEST_SNAPSHOT
```

4. Commit the operation to populate the defaults, as provided by the browser interface. At this point, the operation can be committed, though you will likely need to change the defaults to match the information.

```
delphix database provision defaults *> commit
```

5. Set the name and group for the new VDB.

```
delphix database provision *> set container.name=vexample
delphix database provision *> set container.group="<New Group>"
```

6. Set the base mount point.

```
delphix database provision *> set source.mountBase=/mnt
```

7. Set the source config type to be a single instance non-multitenant Oracle, and set the database name and database unique name. When provisioning from a RAC or single instance non-multitenant oracle source, the default type will match that of the repository selected by the defaults operation.

```
delphix database provision *> set sourceConfig.type=OracleSIConfig
delphix database provision *> set sourceConfig.databaseName=vexample
delphix database provision *> set sourceConfig.uniqueName=vexample123
```

8. Set the instance name and number.

```
delphix database provision *> edit sourceConfig.instance
delphix database provision sourceConfig.instance *> set instanceNumber=1
delphix database provision sourceConfig.instance *> set instanceName=vexample
delphix database provision sourceConfig.instance *> back
```

9. Set the target repository.

```
delphix database provision *> set sourceConfig.repository='/env/opt/oracle'
```

10. Configure the Oracle database parameters. If you are using manually specified parameters, you can set the contents of `source.configParams`. If you want to use a template, you can set `source.configTemplate`.
11. (Optional) Configure customer environment variables.

Setting Environment Pair Values

```
delphix database provision *> edit customEnvVars
delphix database provision source.customEnvVars *> add
delphix database provision source.customEnvVars 0 *> set
type=OracleCustomEnvVarSIPair
delphix database provision source.customEnvVars 0 *> set varName=MYVAR1
delphix database provision source.customEnvVars 0 *> set varValue=Value1
delphix database provision source.customEnvVars 0 *> back
```

```
delphix database provision source.customEnvVars 1 *> set
type=OracleCustomEnvVarSIPair
delphix database provision source.customEnvVars 1 *> set varName=MYVAR2
delphix database provision source.customEnvVars 1 *> set varValue=Value2
delphix database provision source.customEnvVars 1 *> back
```



- **OracleCustomEnvVarSIFile** file should be in the following format.

```
export VARNAME1=VARVALUE1
export VARNAME2=VARVALUE2
```

- When provisioning to a RAC target, use the type as **OracleCustomEnvVarRACPair** or **OracleCustomEnvVarRACFile**.

```
delphix database provision source.customEnvVars 0 *> set
type=OracleCustomEnvVarRACPair
delphix database provision source.customEnvVars 0 *> set
clusterNode=targetnode1
delphix database provision source.customEnvVars 0 *> set
varName=MYVAR1
delphix database provision source.customEnvVars 0 *> set
varValue=Value1
delphix database provision source.customEnvVars 0 *> back
```

12. Commit the result.

```
delphix database provision *> commit
```

CLI Cookbook: Provisioning a SQL Server VDB

This topic describes how to provision a SQL Server VDB using the command line interface.

Prerequisites

You will need the following information:

- The name of the VDB you want to create
- The group in which to create the VDB
- The SQL Server database name for the VDB
- The source dSource or VDB from which you wish to provision
- The semanticLocation, LSN, or timestamp of the point you want to provision from. You can run these commands to get the list of snapshots or timeflow ranges:

```
snapshot list database=dexample
snapshot list timeflow=dexample
snapshot list fromDate="2020-03-01T00:00:00.000Z" toDate="2020-03-04T11:31:27.883Z"
```

- The target host on which you want to create the VDB. You can list the hosts with the `/host list` command.
- The source repository (SQL Server instance on the target host) in which to create the VDB. These can be listed with the `/repository list` command.

Procedure

1. Execute the `database provision` command.

```
delphix> database provision
```

2. Execute the `defaults` command.

```
delphix database provision> defaults
```

3. Set the timeflow point source timeflow and location.

```
delphix database provision defaults *> set type=TimeflowPointSemantic
delphix database provision defaults *> set container=dexample
delphix database provision defaults *> set location=LATEST_SNAPSHOT
```

4. Commit the operation to populate the defaults, as provided by the browser interface. At this point, the operation can be committed, though you will likely need to change the defaults to match the information.

```
delphix database provision defaults *> commit
```

5. Set the name and group for the new VDB.

```
delphix database provision *> set container.name=vexample  
delphix database provision *> set container.group="<New Group>"
```

6. Set the database name for the VDB on the target SQL Server instance.

```
delphix database provision *> set sourceConfig.databaseName=vexample
```

7. Set the target repository.

```
delphix database provision *> set sourceConfig.repository=targetEnv/  
SQLServer2008
```

8. Commit the result.

```
delphix database provision *> commit
```

CLI Cookbook: Provisioning a VDB from a Timeflow Bookmark

This topic describes how to create a Timeflow bookmark and use it to provision a single instance Oracle VDB using the Delphix Engine command-line interface.

You can create Timeflow bookmarks to give a semantically meaningful name to a TimeFlow point (scn, location or timestamp within a Timeflow). You can then use the bookmarks you created to execute the following database operations:

- Provision
- Refresh
- Export
- Test file mappings
- VDB Rewind

Prerequisites

You will need the following information:

- The name of the Timeflow bookmark you want to create
- The name of the VDB you want to create
- The group in which to create the VDB
- The Oracle database name
- The Oracle database unique name
- The Oracle database instance number
- The Oracle database instance name
- The source dSource or VDB from which you wish to provision
- The SCN, or timestamp of the point you want to provision from. You can run these commands to get the list of snapshots or Timeflow ranges:

```
snapshot list database=dexample
timeflow "dexample" timeflowRanges; commit
```

- The base mountpoint on the target server where VDB data should be mounted
- The source repository (oracle install) in which to create the VDB. These can be listed with the `repository list` command.

Creating the Timeflow Bookmark

1. Execute the `timeflow bookmark create` command.

```
delphix> timeflow bookmark create
```

2. Set the timeflow point to be Oracle timeflow point.

```
delphix timeflow bookmark create *> set timeflowPoint.type=OracleTimeflowPoint
```

3. Set the timeflow point timeflow and location

```
delphix timeflow bookmark create *> set timeflowPoint.timeflow=dexample/default
```

```
delphix timeflow bookmark create *> set timeflowPoint.location=1945519455791
```

4. Set the name of the timeflow bookmark

```
delphix timeflow bookmark create *> set name=myTimeFlowBookmark
```

5. Commit the result

```
delphix timeflow bookmark create *> commit
TIMEFLOW_BOOKMARK-1
```

6. Display the list of timeflow bookmarks

```
delphix> timeflow bookmark ls
Objects
NAME                TAG    TIMEFLOW
myTimeFlowBookmark -      dexample/default
Operations
create
```

Provisioning from a TimeFlow Bookmark

1. Execute the `database provision` command.

```
delphix> database provision
```

2. Set `defaults` and provide container (VDB or dSource) that you will be provisioning from

```
delphix database provision > defaults
delphix database provision defaults > set container=<VDB or dSource>
delphix database provision defaults > commit
```

3. Set the `timeflowPointParameters` type to be `TimeflowBookmark`.

```
delphix database provision *> set
timeflowPointParameters.type=TimeflowPointBookmark
```

4. Set the timeflow bookmark.

```
database provision *> set timeflowPointParameters.bookmark=myTimeFlowBookmark
```

5. Set the name and group for the new VDB.

```
delphix database provision *> set container.name=vexample
delphix database provision *> set container.group="Untitled"
```

6. Set the base mountpoint.

```
delphix database provision *> set source.mountBase=/mnt
```

7. Set the source config type to be single instance Oracle, and set the database name and database unique name.

```
delphix database provision *> set sourceConfig.type=OracleSIConfig  
delphix database provision *> set sourceConfig.databaseName=vexample  
delphix database provision *> set sourceConfig.uniqueName=vexample123
```

8. Set the instance name and number.

```
delphix database provision *> edit sourceConfig.instance  
delphix database provision sourceConfig.instance *> set instanceNumber=1  
delphix database provision sourceConfig.instance *> set instanceName=vexample  
delphix database provision sourceConfig.instance *> back
```

9. Set the target repository.

```
delphix database provision *> set sourceConfig.repository=env/'/opt/oracle'
```

10. Commit the result.

```
delphix database provision *> commit
```

CLI Cookbook: Provisioning a Virtual PDB

This topic describes how to provision a virtual pluggable database (vPDB) to a virtual container database (vCDB) using the command-line interface.

⚠ This process applies to Oracle container databases (CDBs) and the pluggable databases (PDBs) found within them that have been linked to Delphix as dSources. As container databases can only be found in Oracle 12.1.0.1 releases and above of Oracle this topic is relevant to only these releases.

Prerequisites

The provisioning of virtual PDBs in Delphix is dependent on the following entities being in place prior to attempting the creation of this type of Delphix dataset.

- A source container database must be linked to the engine.
- The source container database must have pluggable databases within it linked as dSources to Delphix.
- A destination host with an Oracle home of the same release as the source database must be available and already discovered as an Environment in Delphix.

In the example CLI provision detailed below the following databases are in place:

- The source container database is called "CDOMLOSRTB"
- The pluggable database found in "CDOMLOSRTB" that has been linked as a dSource is called "CDOMLOSRTBPDB1"
- The destination virtual container database will be called "cdbvirt"
- Within the Delphix Engine CLI the name of the repository (target Oracle home) on the target system is '/u01/app/oracle/product/12.2.0.1/dbhome_1'
- The destination virtual pluggable database will be called "PDBS3" and run from virtual container database "cdbvirt"
- Delphix creates a temporary container database on the target host which will be used to establish the virtualized copy of the dSource PDB "pborcl". This temporary CDB will be created and running during the provisioning process, it will be destroyed at the end of the provision of the virtual PDB.

Procedure

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
$ ssh admin@YOUR_ENGINE
```

2. Move to the database provisioning command line object.

```
delphix> database provision
```

3. Give the dataset a name.

```
delphix database provision *> set container.name=PDBS3
```

4. Place the new dataset in a Group that appears in the Delphix GUI, in this case, the Targets group.

```
delphix database provision *> set container.group=Targets
```

5. Set the type of provision to perform, for Oracle virtual database (VDB) - VDB/vPDBs, the type will be OracleVirtualPdbSource.

```
delphix database provision *> set source.type=OracleVirtualPdbSource
```

6. Set the destination mount point which Delphix NFS mounts are to be linked to under the virtual PDB. This folder must exist at a file system level on the target host. Do not use single quotes around the mount path.

```
delphix database provision *> set source.mountBase="/mnt/provision"
```

7. Name the vPDB. This is what it will appear as in the destination container database.

```
delphix database provision *> set sourceConfig.databaseName=PDBS3
```

8. Supply the dSource PDBs details. In this example, the provision will use the latest point in time available to the dSource PDB as the point in time from which to provision the vPDB. Setting a different timeflowPointParameters.type would allow you to use points in time other than the latest snapshot or latest point in time if this is what you desire. Using other types is not covered in this example.

```
delphix database provision *> set
timeflowPointParameters.container=CDOMLOSRTBPDB1
```

9. Give the virtual CDB a name:

```
delphix database provision *> set virtualCdb.container.name=cdbvirt
```

10. Place the virtual CDB in a Group that appears in the Delphix GUI, in this case, the target group.

```
delphix database provision *> set virtualCdb.container.group=Targets
```

11. If automatically restarting the vPDB and vCDB is not required after a reboot of the target host, set this to option to false.

```
delphix database provision *> set
virtualCdb.source.allowAutoVDBRestartOnHostReboot=false
```

12. Set the destination mount point which Delphix NFS mounts are to be linked to under the virtual CDB. This folder must exist at a file system level on the target host. Do not use single quotes around the mount path.

```
delphix database provision *> set virtualCdb.source.mountBase="/mnt/provision"
```

13. Set the vCDB configuration type to one of OracleRACConfig or OracleSIConfig. RAC configurations are not covered in this article:

```
delphix database provision *> set virtualCdb.sourceConfig.type=OracleSIConfig
```

14. Set the database name for the virtual CDB:

```
delphix database provision *> set virtualCdb.sourceConfig.databaseName=cdbvirt
```

15. Set the target system environment user that will be used to run the virtual CDB:

```
delphix database provision *> set
virtualCdb.sourceConfig.environmentUser=oracle
```

16. Set the instance name and number for the virtual CDB

```
delphix database provision *> set
virtualCdb.sourceConfig.instance.instanceName=cdbvirt
delphix database provision *> set
virtualCdb.sourceConfig.instance.instanceNumber=1
```

17. Set the repository (Oracle home) on the target system that will be used to run the virtual CDB. In this example the Oracle home is /u01/app/oracle/product/12.2.0.1/dbhome_1:

```
delphix database provision *> set virtualCdb.sourceConfig.repository='/u01/app/
oracle/product/12.2.0.1/dbhome_1'
```

18. Set the database unique name for the virtual CDB:

```
delphix database provision *> set virtualCdb.sourceConfig.uniqueName=cdbvirt
```

19. Set the following two parameters to true or false, depending on whether the vCDB and vPDB should be restarted automatically following a target host reboot:

```
delphix database provision *> set source.allowAutoVDBRestartOnHostReboot=true
delphix database provision *> set
virtualCdb.source.allowAutoVDBRestartOnHostReboot=true
```

20. Check that all the settings you require are in place using the "ls" command.

```
delphix database provision *> ls
Properties
  type: OracleMultitenantProvisionParameters
  container:
    type: OracleDatabaseContainer
    name: PDBS3 (*)
    description: (unset)
    diagnoseNoLoggingFaults: true
    group: Untitled (*)
    performanceMode: DISABLED
    preProvisioningEnabled: false
    sourcingPolicy: (unset)
  credential: (unset)
  masked: (unset)
```

```

maskingJob: (unset)
source:
  type: OracleVirtualPdbSource
  name: (unset)
  allowAutoVDBRestartOnHostReboot: true (*)
  config: (unset)
  customEnvVars: (unset)
  fileMappingRules: (unset)
  logCollectionEnabled: false
  mountBase: /mnt/provision (*)
  operations: (unset)
sourceConfig:
  type: OraclePDBConfig
  cdbConfig: (unset)
  databaseName: cdbvirt (*)
  environmentUser: (unset)
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: (unset)
  services: (unset)
timeflowPointParameters:
  type: TimeflowPointSemantic
  container: CDOMLOSRTBPDB1 (*)
  location: LATEST_POINT
username: (unset)
virtualCdb:
  type: OracleVirtualCdbProvisionParameters (*)
  container:
    type: OracleDatabaseContainer (*)
    name: cdbvirt (*)
    description: (unset)
    diagnoseNoLoggingFaults: true (*)
    group: Untitled (*)
    performanceMode: DISABLED (*)
    preProvisioningEnabled: false (*)
    sourcingPolicy: (unset)
  source:
    type: OracleVirtualCdbSource (*)
    name: (unset)
    allowAutoVDBRestartOnHostReboot: true (*)
    config: (unset)
    configParams: (unset)
    configTemplate: vcdb (*)
    logCollectionEnabled: false (*)
    mountBase: /mnt/provision (*)
  sourceConfig:
    type: OracleSIConfig (*)
    databaseName: cdbvirt (*)
    environmentUser: (unset)
    instance:
      type: OracleInstance (*)

```

```

        instanceName: cdbvirt (*)
        instanceNumber: 1 (*)
        linkingEnabled: true (*)
        nonSysCredentials: (unset)
        nonSysUser: (unset)
        repository: '/u01/app/oracle/product/12.2.0.1/dbhome_1' (*)
        services: (unset)
        tdeKeystorePassword: (unset)
        uniqueName: cdbvirt (*)

```

Operations
defaults

21. Initiate the provision by committing the operation in the CLI.

```

delphix database provision *> commit
PDBS3
Dispatched job JOB-24
DB_PROVISION job started for "Targets/PDBS3".
Starting provision of the virtual database "PDBS3".
Preparing multitenant container database "cdbvirt".
Creating new TimeFlow.
Generating recovery scripts.
Exporting storage.
Mounting filesystems to recover pluggable database on instance "1".
Mounting read-only archive log filesystem for the virtual database instance
"1".
Backing up Oracle spfile.
Mounting virtual database instance.
Disabling flashback on Oracle database.
Renaming Oracle datafiles.
Mounting virtual database instance.
Recovering Oracle pluggable database.
Creating control file.
Creating Oracle online logs.
Processing startup init file.
Configuring initialization and server parameter files.
Mounting virtual database instance.
Performing Oracle resetlogs.
Creating tempfiles.
Renaming readonly datafiles.
Finalizing Oracle pluggable database.
Registering listeners.
Unmounting read-only archive log filesystem for the virtual database
instance "1".
Unmounting filesystems after recovering pluggable database on instance "1".
Plugging in Oracle pluggable database.
Opening Oracle pluggable database.
Setting OMF destination for Oracle pluggable database.
Creating tempfiles.
Online readonly tablespaces.
Enabling Oracle instances.

```

Checking Oracle pluggable database plugin violations.
DB_PROVISION job **for** "Targets/PDBS3" completed successfully.

CLI Cookbook: Provisioning a Virtual PDB in a Target CDB

This topic describes how to provision a virtual pluggable database (vPDB) using the command-line interface.

⚠ This process applies to Oracle container databases (CDBs) and the pluggable databases (PDBs) found within them that have been linked to Delphix as dSources. As container databases can only be found in Oracle **12.1.0.1** releases and above of Oracle this topic is relevant to only these releases.

Prerequisites

Provisioning of virtual PDB's in Delphix is dependent on the following entities being in place prior to attempting the creation of this type of Delphix dataset.

- A source container database must be linked to the engine.
- The source container database must have pluggable databases within it linked as dSources to Delphix.
- A destination container database of the same Oracle release as the source database must exist in the target host and also be linked to Delphix.

In the example CLI provision detailed below the following databases are in place:

- The source container database is called "cdb12"
- The destination container database is called "cdbstage"
- The pluggable database found in "cdb12" that has been linked as a dSource is called "pdborcl"
- The destination virtual pluggable database will be called "pds2" and run from container database "cdbstage"

Delphix creates a temporary container database on the target host which will be used to establish the virtualized copy of the dSource PDB "pdborcl". This temporary CDB will be created and running during the provisioning process, it will be destroyed at the end of the provision of the virtual PDB.

Procedure

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
$ ssh admin@YOUR_ENGINE
```

2. Move to the database provisioning command line object.

```
delphix> database provision
```

3. Set the type of provision to perform, for Oracle virtual database (VDB) - VDB/vPDBs, the type will be OracleVirtualSource.

```
set source.type=OracleVirtualPdbSource
```

4. Set the destination target environment/host through setting the sourceConfig environment user to perform the provision. A destination container database must already be running on this target host.

```
delphix database provision *> set sourceConfig.environmentUser=OEL6SIN1/delphix
```

- Set the destination mount point which Delphix NFS mounts are to be linked to under the virtual PDB. This folder must exist at a file system level on the target host. Do not use single quotes around the mount path.

```
delphix database provision *> set source.mountBase="/mnt/provision"
```

- Set the login details for the provision and Delphix OS user who is to perform the provision.

```
delphix database provision *> set username=delphix
delphix database provision *> set credential.type>PasswordCredential
delphix database provision *> set credential.password=delphix
```

- Give the dataset a name.

```
delphix database provision *> set container.name=PDBS2
```

- Place the new dataset in a Group that appears in the Delphix GUI, in this case, the target group.

```
delphix database provision *> set container.group=Targets
```

- If automatically restarting the VDB is not required after a reboot of the VDB target host, set this to option to false. False is possibly a better option given the container database would need to be running prior to any attempt to pull up a vPDB.

```
delphix database provision *> set source.allowAutoVDBRestartOnHostReboot=false
```

- Supply the destination container database name. This will be where the vPDB will ultimately be placed and run from on the target host.

```
delphix database provision *> set sourceConfig.cdbConfig=cdbstage
```

- Name the vPDB. This is what it will appear as in the destination container database.

```
delphix database provision *> set sourceConfig.databaseName=pdb2
```

- Supply the dSource PDBs details. In this example, the provision will use the latest point in time available to the dSource PDB as the point in time from which to provision the vPDB. Setting a different `timeflowPointParameters.type` would allow you to use points in time other than the latest snapshot or latest point in time if this is what you desire. Using other types is not covered in this example.

```
delphix database provision *> set timeflowPointParameters.container=PDBORCL
```

- Check that all the settings you require are in place using the "ls" command.

```
delphix database provision *> ls
Properties
  type: OracleProvisionParameters
```

```

container:
  type: OracleDatabaseContainer
  name: PDBS2 (*)
  description: (unset)
  diagnoseNoLoggingFaults: true
  group: Targets (*)
  performanceMode: DISABLED
  preProvisioningEnabled: false
  sourcingPolicy: (unset)
credential:
  type: PasswordCredential (*)
  password: ***** (*)
maskingJob: (unset)
newDBID: true (*)
openResetlogs: true
physicalStandby: false
source:
  type: OracleVirtualSource (*)
  name: (unset)
  allowAutoVDBRestartOnHostReboot: false (*)
  archiveLogMode: true
  config: (unset)
  configParams: (unset)
  configTemplate: (unset)
  customEnvVars: (unset)
  fileMappingRules: (unset)
  manualProvisioning: false
  mountBase: /mnt/provision (*)
  nodeListenerList: (unset)
  operations: (unset)
  redoLogGroups: 3
  redoLogSizeInMB: 0
sourceConfig:
  type: OraclePDBConfig
  cdbConfig: cdbstage (*)
  databaseName: pds2 (*)
  environmentUser: (unset)
  linkingEnabled: true
  repository: (unset)
  services: (unset)
timeflowPointParameters:
  type: TimeflowPointSemantic
  container: PDBORCL (*)
  location: LATEST_POINT
  username: delphix (*)

Operations
defaults

```

14. Initiate the provision by committing the operation in the CLI.

```
delphix database provision *> commit
```

PDBS2

Dispatched job JOB-333

DB_PROVISION job started **for** "Targets/PDBS2".

Starting provision of the virtual database "pdbc2".

Preparing multitenant container database "cdbstage".

Creating **new** TimeFlow.

Generating recovery scripts.

Exporting storage.

Mounting filesystems to recover pluggable database on instance "1".

Mounting read-only archive log filesystem **for** the virtual database instance "1".

Recovering Oracle pluggable database.

Mounting filesystems **for** the virtual database instance "1".

Unmounting filesystems after recovering pluggable database on instance "1".

Cleaning up objects created by pluggable database provisioning.

Opening Oracle pluggable database.

DB_PROVISION job **for** "Targets/PDBS2" completed successfully.

CLI Cookbook: Refresh a VDB from a Specific Timepoint or Latest

This topic describes the steps to Refresh a VDB from a specific Timepoint or from Latest.

You can refresh from any point on Timeflow using SCN, location, or timestamp.

Prerequisites

You will need the following information:

- The VDB name
- The TimeFlow location, SCN, or timestamp of the point you want to provision from.

 The default domain user created on Delphix Engines is now **admin** instead of `delphix_admin`. When engines created before 5.3.1 are upgraded to 5.3.1 or later they will retain their old username 'delphix_admin'. To avoid complications Delphix recommends creating users with an admin role and then Disabling `delphix_admin`.

Login to CLI:

```
$ ssh admin@<delphixengine>
```

Determine the TimeFlow

Run:

```
> timeflow "<dSource>" timeflowRanges
> commit
> cd
```

Perform the Refresh from specific Timepoint

```
> database
> select <VDB name>
> refresh
> set timeflowPointParameters.type= (one of TimeflowPointBookmark,
TimeflowPointBookmarkTag, TimeflowPointLocation, TimeflowPointSemantic,
TimeflowPointTimestamp as appropriate)
> set timeflowPointParameters.location= (the location, timestamp, or bookmark you
wish to refresh to)
> set timeflowPointParameters.timeflow= (the timeflow associated with location)
> commit
```

Perform the Refresh from Latest

```
> database
> select <yourdatabase>
> refresh
> set timeflowPointParameters.container= (Parent of VDB)
```

```
> commit
```

CLI Cookbook: Repairing a Timeflow

Prerequisites

- Know the dSource and Group you need to repair from
- Make sure that your retention policy is set correctly so that the ingested logs are within the retention

Procedure

1. Log into the Delphix Engine as an Admin user. Go to Timeflow and then list. Find the name of the Timeflow that needs to be repaired.

```
ssh admin@<yourengine>
delphix > timeflow
delphix timeflow > ls
```

2. List the missing logs for the Timeflow. The maximum number of logs reported is controlled by the value of the pageSize argument; if there are a very large number of missing logs, you may need to increase this value. Note the start and end scn of the missing log.

```
delphix timeflow > cd oracle/log
delphix timeflow oracle log> list timeflow=example missing=true pageSize=1000
```

 If the output from the above command does not list any logs as missing, but there are one or more sequences reported as missing by the **Delphix Admin** browser-based app, please contact Delphix Technical Support for assistance.

3. Stage the missing logs.
 - a. Verify that there is sufficient free space.
 - b. Copy or restore the missing archive logs into an alternative directory on a server the Delphix Engine can access via the network. All files and subdirectories in the configured directory are searched recursively, so staging the archive logs to a location with no other files or folders will ensure the process completes in a timely manner.
 - c. Verify that the user being specified in the next step has permission to read these archive log files in the directory. The user credentials are optional if the host and user credentials have already been added to the Delphix Engine.

```
delphix timeflow oracle log > fetch
delphix timeflow oracle log fetch *> set type=TimeflowLogFetchParameters
delphix timeflow oracle log fetch *> set timeflow=example
delphix timeflow oracle log fetch *> set directory=[directory where you
restored the log file]
delphix timeflow oracle log fetch *> set endLocation=[end SCN of the
sequence]
delphix timeflow oracle log fetch *> set startLocation=[start SCN of the
sequence]
delphix timeflow oracle log fetch *> set host=[hostname or IP of the host
you restored the file to]
```

```
delphix timeflow oracle log fetch *> set username=[a user that can read  
the file]  
delphix timeflow oracle log fetch *> edit credentials  
delphix timeflow oracle log fetch *> set type>PasswordCredential  
delphix timeflow oracle log fetch *> set password=[password for this user]
```

4. Commit the changes.

```
delphix timeflow oracle log fetch *> commit
```

 Only do ONE repair job at a time.

NOTE:

It is possible for there to be more than one Timeflow visible for a given container/source. If that is the case, you can verify the current TimeFlow being used with:

```
delphix > database  
delphix database > select 'example'  
delphix database "example"> ls
```

Look for the `currentTimeflow` value.

CLI Cookbook: Rolling Back a VDB

The following sections provide examples of how to rollback or rewind your VDB.

- [Rolling Back or Rewinding to a Snapshot From a VDB](#)
- [Rolling Back or Rewinding to a Snapshot Using Timeflow](#)
- [Rolling Back or Rewinding to a Timeflow Bookmark](#)

Rolling Back or Rewinding to a Snapshot From a VDB

1. Log into the Delphix Engine.

```
ssh admin@delphix_engine
```

2. List Timeflows for the database that you want to rollback to.

```
de > ls
de > timeflow
de timeflow > ls
```

3. Switch to the VDB you want to rollback.

```
de timeflow > cd /database
de database > ls
de database > select "vdb_example"
```

4. Rollback VDB using the VDB rollback function (note this can also be done in the GUI).

```
de database 'vdb_example' > rollback
de database 'vdb_example' rollback *> set timeflowPointParameters.container=
de database 'vdb_example' rollback *> set timeflowPointParameters.location=
de database 'vdb_example' rollback *> commit
```

Rolling Back or Rewinding to a Snapshot Using Timeflow

1. Log into the Delphix Engine.

```
ssh admin@delphix_engine
```

2. List Timeflows for the database that you want to rollback to.

```
de > ls
de > timeflow
de timeflow > ls
```

3. Switch to the VDB you want to rollback.

```
de timeflow > cd /database
de database > ls
de database > select "vdb_example"
```

4. Use the switchTimeflow operation.

```
de database 'vdb_example' > switchTimeflow
de database 'vdb_example' switchTimeflow *> set timeflow=<different timeflow>
de database 'vdb_example' switchTimeflow *> commit
```

Rolling Back or Rewinding to a Timeflow Bookmark

Requirements: Know the Timeflow bookmark that you want to use.

1. Log into the Delphix Engine.

```
ssh admin@<yourengine>
```

2. Retrieve database and Timeflow information that you would like to rewind/rollback to.

```
delphix > ls
delphix database > select "dexample"
delphix database "dexample" > get currentTimeflow
```

3. Rollback/Rewind VDB.

```
delphix database "dexample" > rollback
delphix database "dexample" rollback *> ls
delphix database "dexample" rollback *> set
timeflowPointParameters.type=TimeflowPointBookmark
delphix database "dexample" rollback *> set timeflowPointParameters.bookmark="bookmar
k example"
delphix database "dexample" rollback *> commit
```

CLI Cookbook: Rolling Forward a VDB

This topic describes how to roll forward a virtual database after it has been rewound, as described in [Provisioning and Managing Virtual Databases](#).

Once a VDB has rewound to a specific TimeFlow point, the snapshots of its previous states are still available in Delphix Engine storage and are accessed via the command-line interface to restore those previous states. This is referred to as "rolling forward" a VDB.

Procedure

1. Use the `ls` command to find the VDB you want to roll forward.

In this example, the TimeFlows and their associated containers are listed. The VDB called `PVDB` will be the one to roll forward.

```
delphix timeflow> ls
Objects
NAME                                CONTAINER  PARENTPOINT.TIMEFLOW
PARENTPOINT.LOCATION  PARENTPOINT.TIMESTAMP
hrprod/default        hrprod     -
-
erpprod/default       erpprod    -
-
'DB_PROVISION@2013-11-25T17:37:06'  PVDB       erpprod/default
657925
'DB_ROLLBACK@2013-11-25T18:24:16'  PVDB
'DB_PROVISION@2013-11-25T17:37:06'  678552
```

2. Use the `Select` command to select the database.

```
delphix database> select PVDB
```

3. Use the `rollback` command to roll forward the VDB.

```
delphix database "PVDB"> rollback
```

4. Use the `ls` command to display options for selecting TimeFlow parameters.

```
delphix database "PVDB" rollback *> ls
Properties
  type: OracleRollbackParameters
  credential: (unset)
  timeflowPointParameters:
    type: TimeflowPointSemantic
    container: (required)
    location: LATEST_POINT
  username: (unset)
```

5. Because this VDB was rolled back, two TimeFlows now exist for it. To rollback the VDB and roll it forward, select the original TimeFlow, because the original snapshots are associated with that TimeFlow.

```
delphix database "PVDB" rollback *> set
timeflowPointParameters.type=TimeflowPointLocation
delphix database "PVDB" rollback *> set timeflowPointParameters.timeflow='DB_PROVISION@2013-11-25T17:37:06'
```

6. Use the `ls` command to view the parameter options for the TimeFlow you selected.

```
delphix database "PVDB" rollback *> ls
Properties
  type: OracleRollbackParameters
  credential: (unset)
  timeflowPointParameters:
    type: TimeflowPointLocation (*)
    location: LATEST_POINT
    timeflow: 'DB_PROVISION@2013-11-25T17:37:06' (*)
  username: (unset)
```

7. Set the TimeFlow location to rollback the VDB to a particular Oracle SCN.

```
delphix database "PVDB" rollback *> set timeflowPointParameters.location=678994
```

8. Use the `ls` command to review all the options you selected before executing the commit.

```
delphix database "PVDB" rollback *> ls
Properties
  type: OracleRollbackParameters
  credential: (unset)
  timeflowPointParameters:
    type: TimeflowPointLocation (*)
    location: 678994 (*)
    timeflow: 'DB_PROVISION@2013-11-25T17:37:06' (*)
  username: (unset)
```

9. `Commit` the changes.

```
delphix database "PVDB" rollback *> commit
Dispatched job JOB-369
DB_ROLLBACK job started for "ERP/PVDB".
Starting provision of the virtual database "PVDB".
Creating new TimeFlow.
Generating recovery scripts.
Exporting storage.
Validating user environment settings on target host.
Mounting filesystems for the virtual database instance "1".
```

```
Mounting read-only archive log filesystem for the virtual database instance
"1".
Running user-specified pre-provisioning script.
Recovering Oracle database.
Opening the virtual database "PVDB".
Opening Oracle database.
Oracle recovery was successful.
Unmounting read-only archive log filesystem for the virtual database
instance "1".
Running user-specified post-provisioning script.
The virtual database "PVDB" was successfully provisioned.
Starting snapshot of virtual database.
Processing database files of virtual database.
Creating snapshot of virtual database.
Finalizing snapshot of virtual database.
Virtual database "PVDB" snapshot successful.
DB_ROLLBACK job for "ERP/PVDB" completed successfully.
```

CLI Cookbook: Taking a Snapshot

This article is to document how to take a Snapshot outside of the normal snapshot policy time using the CLI, you can also do this in the GUI using the camera icon. A Snapshot of a VDB is similar to bookmarking a point in time in the life of the VDB.

Procedure:

1. ssh into the Delphix Engine using delphix_admin credentials.
2. Go into databases and select the VDB or dSource you would like to take a Snapshot of .

```
ssh admin@engine
delphix > database
delphix database > select vdb_test
```

3. You are now going to sync and commit the operation .

```
delphix database "vdb_test" > sync
delphix database "vdb_test" sync *> commit
```

4. You can verify the snapshot by going to snapshots and listing them

```
delphix database "vdb_test" > /snapshot
delphix snapshot > ls
```

CLI Cookbook: V2P: Virtual to Physical of a Single Instance Non-multitenant Oracle Database

This topic describes how to provision a physical single instance non-multitenant Oracle database using the Delphix Engine command-line interface.

Prerequisites

You will need the following information:

- The instance name, instance number, and unique name of the Oracle database you wish to create
- The source dSource or VDB from which you wish to provision.
- The semanticLocation, SCN, or timestamp of the point you want to provision from. You can run these commands to get the list of snapshots or Timeflow ranges:

```
snapshot list database=dexample
timeflow "dexample" timeflowRanges; commit
```

- The layout of the filesystems on the target server where data should be exported.
- The source repository (oracle install) in which to create the VDB. These can be listed with the `repository list` command.

Procedure

1. Execute the `database export` command.

```
delphix> database export
```

2. Set the Timeflow point type, source container, and location.

```
delphix database export *> set
timeflowPointParameters.type=TimeflowPointSemantic
delphix database export *> set timeflowPointParameters.container=dexample
delphix database export *> set timeflowPointParameters.location=LATEST_SNAPSHOT
```

3. Edit the sourceConfig configuration, specifying the parameters for the database created via V2P.

```
delphix database export *> edit sourceConfig
delphix database export sourceConfig *> set type=OracleSIConfig
delphix database export sourceConfig *> edit instance
delphix database export sourceConfig instance *> ls
delphix database export sourceConfig instance *> set instanceName=v2p_db
delphix database export sourceConfig instance *> set instanceNumber=1
delphix database export sourceConfig instance *> back
delphix database export sourceConfig *> set repository=tserver/'/u01/app/
ora11204/product/11.2.0/dbhome_1'
delphix database export sourceConfig *> set uniqueName=v2p_db
delphix database export sourceConfig *> set databaseName=v2p_db
delphix database export sourceConfig *> back
```

4. Set the destination locations.

```
delphix database export *> set filesystemLayout.targetDirectory=//u01/app/  
oracle/oradata/v2p_db  
delphix database export *> set filesystemLayout.archiveDirectory=archive  
delphix database export *> set filesystemLayout.dataDirectory=datafiles  
delphix database export *> set filesystemLayout.externalDirectory=external  
delphix database export *> set filesystemLayout.scriptDirectory=script  
delphix database export *> set filesystemLayout.tempDirectory=temp
```

5. Commit the configuration to execute the job.

```
delphix database export *> commit
```

CLI Cookbook: V2P (Virtual to Physical) of a Single Instance Oracle Database with Datafiles on Separate File Systems

This topic describes how to provision a physical single-instance Oracle database, with datafiles on separate file systems, using the Delphix Engine command-line interface.

By default, all of the customizable directories - `data`, `archive`, `temp`, `external`, and `script` will be placed under the `target` directory, with additional config files placed in the `target` directory itself. This provides no mechanism to place datafiles on separate file systems (i.e. `/data1` and `/data2`, which have the root directory as a common point). In order to allow this, the `useAbsolutePathForDataFiles` parameter must be specified.

Prerequisites

You will need the following information:

- The instance name, instance number, and unique name of the Oracle database to create.
- The source dSource or VDB from which to provision.
- The semanticLocation, SCN, or timestamp of the point of which to provision from. Run these commands to get the list of snapshots or Timeflow ranges:

```
snapshot list database=dexample
timeflow "dexample" timeflowRanges; commit
```

- The layout of the filesystems on the target server where data should be exported.
- The source repository (oracle install) in which to create the VDB. These can be listed with the `repository list` command.

Procedure

1. Execute the `database export` command.

```
delphix> database export
```

2. Set the Timeflow point type, source container, and location.

```
delphix database export *> set
timeflowPointParameters.type=TimeflowPointSemanticdelphix database export *>
set timeflowPointParameters.container=dexample
delphix database export *> set
timeflowPointParameters.location=LATEST_SNAPSHOT
```

3. Edit the sourceConfig configuration, specifying the parameters for the database created via V2P.

```
delphix database export *> edit sourceConfig
delphix database export sourceConfig *> set type=OracleSIConfig delphix
database export sourceConfig *> edit instance
delphix database export sourceConfig instance *> ls
delphix database export sourceConfig instance *> set instanceName=v2p_db
```

```
delphix database export sourceConfig instance *> set instanceNumber=1delphix
database export sourceConfig instance *> back
delphix database export sourceConfig *> set repository=tserver/'/u01/app/
ora11204/product/11.2.0/dbhome_1' delphix database export sourceConfig *> set
uniqueName=v2p_db
delphix database export sourceConfig *> set databaseName=v2p_db
delphix database export sourceConfig *> back
```

4. Set the destination locations, including `setAbsolutePathForDataFiles` to be true.

```
delphix database export *> set filesystemLayout.targetDirectory=/v2p_db
delphix database export *> set filesystemLayout.archiveDirectory=archive
delphix database export *> set filesystemLayout.dataDirectory=datafiles
delphix database export *> set filesystemLayout.externalDirectory=external
delphix database export *> set filesystemLayout.scriptDirectory=script
delphix database export *> set filesystemLayout.tempDirectory=temp
delphix database export *> set
filesystemLayout.useAbsolutePathForDataFiles=true
```

5. Specify the file mapping rules for each datafile and the location to which it should be exported. Each entry should be separated by a newline, and the entire mapping within double quotes.

```
set fileMappingRules="/u01/app/oracle/oradata/datafile1.dbf:/data1/
datafile1.dbf
/u01/app/oracle/oradata/datafile2.dbf:/data1/datafile2.dbf
/u01/app/oracle/oradata/datafile3.dbf:/data2/datafile3.dbf
/u01/app/oracle/oradata/datafile4.dbf:/data2/datafile4.dbf"
```

6. Commit the configuration to execute the job.

```
delphix database export *> commit
```

CLI Cookbook: V2P Virtual to Physical on SQL Server

This topic describes how to export a physical single instance SQL Server database from a VDB using the Delphix Engine command-line interface.

Prerequisites

You will need the following information:

- The VDB name which you wish to export.
- The database name of the SQL Server database you wish to create
- The layout of the filesystems on the target server where data should be exported.
- The target repository in which to create the physical database. Repositories can be listed with the `repository list` command.

Procedure

1. Execute the command `database export` .

```
delphix> database export
```

2. Specify the `set type=MSSqlExportParameters`

```
delphix database export *> set type=MSSqlExportParameters
```

3. Set the VDB you wish to export

```
delphix database export *> edit timeflowPointParameters
delphix database export timeflowPointParameters *> set container=vdb
delphix database export timeflowPointParameters *> back
```

4. Edit the `sourceConfig` configuration, specifying the parameters for the database created via V2P.

```
delphix database export *> edit sourceConfig
delphix database export sourceConfig *> set type=MSSqlSIConfig
delphix database export sourceConfig *> set databaseName=v2p_db
delphix database export sourceConfig *> set repository=win2012/SQL2016
delphix database export sourceConfig *> back
```

5. Set the destination locations.

```
delphix database export *> edit filesystemLayout
delphix database export filesystemLayout *> set targetDirectory=C:\temp
delphix database export filesystemLayout *> set archiveDirectory=archive
delphix database export filesystemLayout *> set dataDirectory=datafiles
delphix database export filesystemLayout *> set externalDirectory=external
delphix database export filesystemLayout *> set scriptDirectory=script
delphix database export filesystemLayout *> set tempDirectory=temp
```

6. Commit the configuration to execute the job.

```
delphix database export *> commit
```

CLI Cookbook: VDB status

It is possible to get a virtual database (VDB) status from the CLI.

1. Log into the CLI as any user that has privileges on the VDB.

```
ssh admin@yourengine
```

2. From source go to the VDB you want to get a status on.

```
delphix > source
delphix source > ls
delphix source > select <yourvdb>
```

3. Run the get runtime command to see all information or just get runtime.status for if the VDB is running.

```
delphix source 'vdb' > get runtime
```

or

```
delphix source 'vdb' > get runtime.status
```

4. If you would like to see more than one VDBs status you can also do the following:

```
delphix > source
delphix source > list display=name,runtime.status
```

CLI Cookbook: Provisioning a Virtual PDB from a Non-Multitenant Source Database

Delphix supports provisioning a vPDB from a non-multitenant source database. This feature is only available through the API or command-line interface.

This topic describes how to provision a virtual pluggable database (vPDB) from a non-multitenant source database using the command-line interface.

- ⚠** This feature has the following restrictions:
1. Transparent Data Encryption (TDE) is not supported.
 2. The provision point must correspond to a snapshot. Provisioning from a point in time between snapshots is not supported.
 3. The target CDB must be a physical CDB. Virtual CDB targets are not supported.

Prerequisites

Provisioning a vPDB from a non-multitenant source has the following environment requirements:

- Source host with a non-multitenant Oracle 11g or newer source database.
- VDB target host for provisioning a virtual non-multitenant VDB from the source database.
- CDB target host with a running Oracle target version CDB. The target CDB will be automatically linked if it is not already linked.

The target CDB can be a newer Oracle version than the source database (for example, the source is 12.2 and target is 19c). When an upgrade is also required, there are two options for upgrading the database:

- Upgrade Option 1: After provisioning the VDB. This option requires the ability to upgrade to the Oracle target version on the VDB target host.
- Upgrade Option 2: After plugging into the Linked CDB target database.

There are three scripts used during this procedure:

1. **Pre-snapshot Hook on VDB**: This hook will open the database in "read only" mode and issue a call to the `dbms_pdb.describe` procedure to generate an XML file describing the VDB.
2. **Post-snapshot Hook on VDB**: This hook will return the VDB to "read write" mode.
3. **Non-CDB to PDB Script**: This script will run as a hook present on the Linked CDB target host. This should call into the Oracle script `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` to convert the VDB into a PDB. This script should also upgrade the vPDB if doing Upgrade Option 2.

Workflow

1. Link the non-multitenant Oracle 11g or newer source database as a dSource within Delphix.
2. Provision a non-multitenant Oracle VDB from the dSource onto the VDB target host. This will be referred to as the **Golden VDB**.
3. (If using Upgrade Option 1) Upgrade the **Golden VDB** to the Oracle target version: manually upgrade the database and point it to the new Oracle home. This step is only necessary if the source and target Oracle versions are not the same and the data files will not be upgraded when they are converted below.
4. Create a **Pre-snapshot hook** on the **Golden VDB** to open the database in read only mode and issue the `dbms_pdb.describe` procedure call to create an XML file called `delphix_plugin.xml`. The XML file will be used to plug the **Golden VDB** data files into the target CDB. The **Golden VDB** must be open read only during the subsequent snapshot so that the VDB data files do not require recovery when plugging them into the target CDB.

5. (Optional) Create a [Post-snapshot hook](#) on the **Golden VDB** to remove the database from read only mode. The **Golden VDB** can also remain read only.
6. Take a snapshot of the **Golden VDB**.
7. Create a [PDB conversion script](#) named `dx-post-plug-hook.sh` in the root of the Delphix toolkit directory of the Linked CDB target host. The name of the vPDB being provisioned/converted will be supplied by Delphix as the first parameter to the script when it invokes the script. The VDB data files will already be plugged into the Linked CDB target at the time the script is invoked. The script should do the following:
 - a. (If using Upgrade Option 2) [Upgrade](#) the vPDB data files prior to the conversion.
 - b. Call into `$ORACLE_HOME/rdbms/admin/noncdb_to_pdb.sql` and perform any customizations for the multitenant conversion.
8. Select a **snapshot** (point-in-time not supported) on the **Golden VDB** that has the `delphix_plugin.xml` file and provision a virtual PDB to the Linked CDB target. Virtual CDB targets are not supported.
Note: This step can be executed via the **API / CLI only**, and will not be allowed via the Delphix UI.

CLI Procedure to Provision a vPDB from a VDB

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
$ ssh admin@YOUR_ENGINE
```

2. Move to the database provisioning command line object.

```
delphix> database provision
```

3. Set the parameter type to `OracleMultitenantProvisionParameters`.

```
set type=OracleMultitenantProvisionParameters
```

4. Set the login details for the provision and Delphix OS user who is to perform the provision.

```
delphix database provision *> set username=delphix
delphix database provision *> set credential.type>PasswordCredential
delphix database provision *> set credential.password=delphix
```

5. Give the dataset a name.

```
delphix database provision *> set container.name=vpdb
```

6. Place the new dataset in a Group that appears in the Delphix GUI, in this case, the Targets group.

```
delphix database provision *> set container.group=Targets
```

7. Set the destination mount point which Delphix NFS mounts are to be linked to under the virtual PDB. This folder must exist at a file system level on the Linked CDB target host. Do not use single quotes around the mount path.

```
delphix database provision *> set source.mountBase="/mnt/provision"
```

- If automatically restarting the vPDB is not required after a reboot of the Linked CDB target host, set this to option to false. False is possibly a better option given the container database would need to be running prior to any attempt to pull up a vPDB.

```
delphix database provision *> set source.allowAutoVDBRestartOnHostReboot=false
```

- Supply the destination container database name. The container database should already be discovered. This will be where the vPDB will ultimately be placed.

```
delphix database provision *> set sourceConfig.cdbConfig=CDBSTAGE
```

- Name the vPDB. This is what it will appear as in the destination container database.

```
delphix database provision *> set sourceConfig.databaseName=vpdb
```

- Supply the source **Golden VDB** details. In this example, the provision will use the latest snapshot available from the **Golden VDB** as the point in time from which to provision the vPDB. A specific snapshot can also be picked, but an arbitrary point in time is not supported.

```
delphix database provision *> set
timeflowPointParameters.type=TimeflowPointSemantic
delphix database provision *> set timeflowPointParameters.container=gold_vdb
delphix database provision *> set
timeflowPointParameters.location=LATEST_SNAPSHOT
```

- Check that all the settings you require are in place using the "ls" command.

```
delphix database provision *> ls
Properties
  type: OracleMultitenantProvisionParameters
  container:
    type: OracleDatabaseContainer
    name: vpdb (*)
    description: (unset)
    diagnoseNoLoggingFaults: true
    group: Targets (*)
    performanceMode: DISABLED
    preProvisioningEnabled: false
    sourcingPolicy: (unset)
  credential:
    type: PasswordCredential (*)
    password: ***** (*)
  masked: (unset)
  maskingJob: (unset)
  source:
    type: OracleVirtualPdbSource (*)
```

```

name: (unset)
allowAutoVDBRestartOnHostReboot: false (*)
config: (unset)
customEnvVars: (unset)
fileMappingRules: (unset)
LogCollectionEnabled: false
mountBase: /mnt/provision (*)
operations: (unset)
parentTdeKeystorePassword: (unset)
parentTdeKeystorePath: (unset)
tdeExportedKeyFileSecret: (unset)
sourceConfig:
  type: OraclePDBConfig
  cdbConfig: CDBSTAGE (*)
  databaseName: vpdb (*)
  environmentUser: (unset)
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: (unset)
  services: (unset)
timeflowPointParameters:
  type: TimeflowPointSemantic
  container: gold_vdb (*)
  location: LATEST_SNAPSHOT (*)
username: delphix (*)
VirtualCdb: (unset)

```

Operations
defaults

13. Initiate the provision by committing the operation in the CLI.

```

delphix database provision *> commit
vpdb
Dispatched job JOB-333
DB_PROVISION job started for "Targets/vpdb".
Starting provision of virtual PDB database "vpdb" converted from a single tenant
database.
Preparing multitenant container database "CDBSTAGE".
Creating new TimeFlow.
Generating recovery scripts.
Exporting storage.
Preparing XML manifest file prior to plugin.
Plugging in Oracle pluggable database.
Running user-defined post plug hook.
Opening Oracle pluggable database.
Setting OMF destination for Oracle pluggable database.
Creating PDB tempfiles.
Checking Oracle pluggable database plugin violations.
DB_PROVISION job for "Targets/vpdb" completed successfully.

```

- ✔ To refresh the data in the vPDB from production, first, refresh the **Golden VDB** from the dSource, then refresh the vPDB from the new snapshot in the **Golden VDB**.

- ⓘ There are some workflow customizations required for RAC databases:
1. The PDB conversion script must be in the root of the Delphix toolkit directory for all the target CDB RAC instances.
 2. The **Golden VDB** Pre-Snapshot hook, as provided below, will not work in a clustered (RAC) environment with more than one active instance because it only shuts down the local instance. `dbms_pdb.describe` will not execute while an instance is open read-write. The workarounds are:
 - a. Provision the **Golden VDB** as single-instance, either by provisioning to a non-RAC target or by provisioning to a RAC target with only one active instance. The sample hook will work in this case.
 - b. Write a customized pre-snapshot hook that shuts down all instances, restarts only one instance in read-only mode, and runs `dbms_pdb.describe`.
 - c. Manually perform the actions of the hook: shutdown the **Golden VDB**, restart one of the instances in read-only mode and then run `dbms_pdb.describe`.

Sample Scripts

Golden VDB Pre-Snapshot Hook

Restarts the source VDB in read only mode and runs `dbms_pdb.describe` to generate an XML file describing the VDB. The XML file will be used to plug the VDB into the Linked CDB target. The target for the XML file must be `$DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/delphix_plugin.xml`.

```
#!/bin/sh

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/presnapshot.out
replace
shutdown immediate
startup mount
alter database open read only;
exec dbms_pdb.describe(pdb_descr_file=>'$DELPHIX_MOUNT_PATH/
$DELPHIX_DATABASE_UNIQUE_NAME/datafile/delphix_plugin.xml');
exit;
EOF
```

Golden VDB Post-Snapshot Hook

This is only necessary if the VDB should not be left in read-only mode after the snapshot.

```
#!/bin/sh

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $DELPHIX_MOUNT_PATH/$DELPHIX_DATABASE_UNIQUE_NAME/datafile/postsnapshot.out
replace
  shutdown immediate
  startup
  exit;
EOF
```

PDB Conversion Script

The script should be named `dx-post-plug-hook.sh` and reside in the root of the Delphix toolkit directory of the Linked CDB target host. Delphix will supply the name of the PDB being provisioned/converted as the first parameter.

The VDB datafiles will have already been plugged into the target CDB at the time the script is invoked and the virtual PDB will be in the mounted (not open) state. The PDB conversion script should return with the virtual PDB in either the mounted or open (not restricted) state. Delphix does not enforce a time-out for the script.

```
#!/bin/sh

DELPHIX_PDB_NAME=$1
SCRIPT_DIR="$( cd "$( dirname "$0" )" && pwd )"
CONVERT_LOGFILE=$SCRIPT_DIR/$DELPHIX_PDB_NAME-pdbconvert.log

sqlplus "/ AS SYSDBA" <<-EOF
  whenever sqlerror exit 2;
  spool $CONVERT_LOGFILE replace
  alter session set container=$DELPHIX_PDB_NAME;
  @?/rdbms/admin/noncdb_to_pdb.sql
  exit;
EOF
```

PDB Upgrade Script

The following script will upgrade the vPDB. Use a wrapper that runs both this script and the prior conversion script (or combine the two in a single script) if doing both an upgrade and a conversion.

```
#!/bin/sh

DELPHIX_PDB_NAME=$1
SCRIPT_DIR="$( cd "$( dirname "$0" )" && pwd )"
UPGRADE_LOGFILE=$SCRIPT_DIR/$DELPHIX_PDB_NAME-dx-post-plug-upgrade.log
UPGRADE_LOGDIR=$SCRIPT_DIR/$DELPHIX_PDB_NAME-upgrade

mkdir $UPGRADE_LOGDIR
cd $ORACLE_HOME/rdbms/admin
switches="-c '$DELPHIX_PDB_NAME' -l $UPGRADE_LOGDIR"
$ORACLE_HOME/perl/bin/perl catctl.pl $switches catupgrd.sql &>> $UPGRADE_LOGFILE
```

CLI Cookbook: Migrating a Virtual PDB and a Virtual CDB

Delphix supports migrating from a vPDB in a virtual CDB, to another linked or virtual CDB. This feature is only available through the command-line interface.

Note

Attempts to perform this migration via GUI will fail with the following error:

Error

The operation could not be completed because the virtual database “vPDB1” is disabled.

Error Code

exception.db.genericvdb.disabled

This topic describes how to migrate a virtual pluggable database (vPDB) in a virtual CDB to another virtual CDB using the command-line interface.

Pre-requisites

You should already set up and have Delphix discover a container database in the same environment as the vPDB currently is or from an environment to which the vPDB will be migrated to.

Procedure

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
ssh admin@YOUR_ENGINE
```

2. Disable the vPDB.

```
delphix> /source
delphix source> select myPDB
delphix source 'myPDB'> disable
delphix source 'myPDB' disable *> commit
  Dispatched job JOB-35
  SOURCE_DISABLE job started for "myPDB".
  Disabling virtual database "myPDB".
  Unexporting storage containers.
  Virtual database "myPDB" disabled.
  SOURCE_DISABLE job for "myPDB" completed successfully.
delphix source 'myPDB'>
```

3. Change the environment user and repository (ORACLE_HOME) for the vPDB to the appropriate user from the new host.

```
delphix> /sourceconfig
delphix sourceconfig> select myPDB
delphix sourceconfig 'myPDB'> update
delphix sourceconfig 'myPDB' update *> set repository='/u01/app/oracle/product/
19.0.0.0/dbhome_1'
delphix sourceconfig 'myPDB' update *> set environmentUser='/oracle'
delphix sourceconfig 'myPDB' update *> commit
```

```

Dispatched job JOB-38
SOURCE_CONFIG_UPDATE job started for "myPDB".
SOURCE_CONFIG_UPDATE job for "myPDB" completed successfully.

```

4. Change the environment user and repository (ORACLE_HOME) for the vCDB to the appropriate user from the new host. This must match the settings in step 3.

```

delphix sourceconfig 'myPDB'> /sourceconfig/
delphix sourceconfig> select myCDB
delphix sourceconfig 'myCDB'> update
delphix sourceconfig 'myCDB' update *> set repository='/u01/app/oracle/product/
19.0.0.0/dbhome_1'
delphix sourceconfig 'myCDB' update *> set environmentUser='/oracle'
delphix sourceconfig 'myCDB' update *> commit
Dispatched job JOB-39
SOURCE_CONFIG_UPDATE job started for "myCDB".
SOURCE_CONFIG_UPDATE job for "myCDB" completed successfully.

```

5. Enable the vPDB.

```

delphix sourceconfig 'myCDB'> /source
delphix source> select myPDB
delphix source 'myPDB'> enable
delphix source 'myPDB' enable *> commit
Dispatched job JOB-40
SOURCE_ENABLE job started for "myPDB".
Enabling dataset "myPDB".
Exporting storage containers from the Delphix Engine.
Mounting datasets.
Mounting filesystems for the virtual database instance "1".
Starting virtual database.
Starting instance 1 on virtual database "myCDB".
Virtual database "myCDB" was successfully started.
Starting virtual database.
Starting instance 1 on virtual database "myPDB".
Plugging in Oracle pluggable database.
Opening Oracle pluggable database.
Setting OMF destination for Oracle pluggable database.
Creating PDB tempfiles.
Checking Oracle pluggable database plugin violations.
Virtual database "myPDB" was successfully started.
Dataset "myPDB" enabled.
SOURCE_ENABLE job for "myPDB" completed successfully.

```

CLI Cookbook: Migrating an Oracle RAC Virtual PDB and a Virtual CDB

Delphix supports migrating from a RAC vPDB in a virtual CDB, to another RAC linked CDB or virtual CDB. This feature is only available through the command-line interface.

Note

Attempts to perform this migration via GUI will fail with the following error:

Error

The operation could not be completed because the virtual database “vPDB1” is disabled.

Error Code

exception.db.genericvdb.disabled

This topic describes how to migrate a RAC virtual pluggable database (vPDB) in a virtual CDB to another RAC virtual CDB using the command-line interface.

Pre-requisites

You should already set up and have Delphix discover a container database in the same environment as the vPDB currently is or from an environment to which the vPDB will be migrated to.

Procedure

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
ssh admin@YOUR_ENGINE
```

2. Disable the vPDB.

```
delphix> /source
delphix source> select racPDB
delphix source 'racPDB'> disable
delphix source 'racPDB' disable *> commit
  Dispatched job JOB-74
  SOURCE_DISABLE job started for "racPDB".
  Disabling virtual database "racPDB".
  Unexporting storage containers.
  Virtual database "racPDB" disabled.
  SOURCE_DISABLE job for "racPDB" completed successfully.
```

3. Change the environment user and repository (ORACLE_HOME) for the vPDB to the appropriate user from the new host.

```
delphix source 'racPDB'> /sourceconfig/
delphix sourceconfig> select racPDB
delphix sourceconfig 'racPDB'> update
delphix sourceconfig 'racPDB' update *> set repository='/u01/app/oracle/
product/12.1.0.2/dbhome_1'
delphix sourceconfig 'racPDB' update *> set environmentUser=' /oracle'
delphix sourceconfig 'racPDB' update *> commit
  Dispatched job JOB-76
```

```
SOURCE_CONFIG_UPDATE job started for "racPDB".
SOURCE_CONFIG_UPDATE job for "racPDB" completed successfully.
```

4. Change the environment user and repository (ORACLE_HOME) for the vCDB to the appropriate user from the new host. This must match the settings in step 3. The instance configuration must also be configured for the new hosts.

```
delphix sourceconfig 'racPDB'> /sourceconfig/
delphix sourceconfig> select racCDB
delphix sourceconfig 'racCDB'> update
delphix sourceconfig 'racCDB' update *> set repository='/u01/app/oracle/
product/12.1.0.2/dbhome_1'
delphix sourceconfig 'racCDB' update *> set environmentUser=' /oracle'
delphix sourceconfig 'racCDB' update *> edit instances
delphix sourceconfig 'racCDB' update instances *> ls
Properties
  0:
    type: OracleRACInstance
    instanceName: racCDB1
    instanceNumber: 1
    node: mwtestx1
  1:
    type: OracleRACInstance
    instanceName: racCDB2
    instanceNumber: 2
    node: mwtestx2
```

Use the "add" command to add an element to **this** array.

```
delphix sourceconfig 'racCDB' update instances *> set 0.node=mwtestz1
delphix sourceconfig 'racCDB' update instances *> set 1.node=mwtestz2
delphix sourceconfig 'racCDB' update instances *> commit
Dispatched job JOB-77
SOURCE_CONFIG_UPDATE job started for "racCDB".
SOURCE_CONFIG_UPDATE job for "racCDB" completed successfully.
```

5. Enable the vPDB.

```
delphix sourceconfig 'racCDB'> /source
delphix source> select racPDB
delphix source 'racPDB'> enable
delphix source 'racPDB' enable *> commit
Dispatched job JOB-78
SOURCE_ENABLE job started for "racPDB".
Enabling dataset "racPDB".
Exporting storage containers from the Delphix Engine.
Mounting datasets.
Mounting filesystems for the virtual database instance "2".
Starting virtual database.
Starting instance 2 on virtual database "racCDB".
Virtual database "racCDB" was successfully started.
Starting virtual database.
```

```
Starting instance 2 on virtual database "racPDB".
Plugging in Oracle pluggable database.
Opening Oracle pluggable database.
Setting OMF destination for Oracle pluggable database.
Creating PDB tempfiles.
Checking Oracle pluggable database plugin violations.
Virtual database "racPDB" was successfully started.
Mounting datasets.
Mounting filesystems for the virtual database instance "1".
Starting virtual database.
Starting instance 1 on virtual database "racCDB".
Virtual database "racCDB" was successfully started.
Starting virtual database.
Starting instance 1 on virtual database "racPDB".
Opening Oracle database.
Checking Oracle pluggable database plugin violations.
Virtual database "racPDB" was successfully started.
Dataset "racPDB" enabled.
SOURCE_ENABLE job for "racPDB" completed successfully.
```

CLI Cookbook: Locating and Updating the Value of tdeKeyIdentifier

This topic describes how to manage the tdeKeyIdentifier field that is associated with the vPDB source object using the command-line interface.

 This process is currently supported only via CLI.

Procedure

1. Log into the Delphix command-line interface using the admin user or a user with admin privileges.

```
$ ssh admin@YOUR_ENGINE
```

2. Move to the database.

```
delphix> source
delphix source> "VCDO_1JL"
```

3. Viewing all the settings using the "ls" command.

```
delphix source "VCDO_1JL" *> ls
Properties
  type: OracleVirtualPdbSource
  name: VCDO_1JL
  allowAutoVDBRestartOnHostReboot: false
  archiveLogMode: true
  config: VCDO_1JL
  configParams:
    _cdb_disable_pdb_limit: TRUE
    audit_file_dest: '/u01/app/oracle/admin/CDOMLOSR197/adump'
    audit_trail: 'DB'
    compatible: '19.0.0'
    core_dump_dest: '/u01/app/oracle/diag/rdbms/cdomlosr197/CDOMLOSR197/
cdump'
    diagnostic_dest: '/u01/app/oracle'
    dispatchers: '(PROTOCOL=TCP) (SERVICE=CDOMLOSRC1DXDB)'
    enable_pluggable_database: TRUE
    log_archive_format: '%t_%s_%r.dbf'
    max_pdbs: 4098
    memory_max_target: 1342177280
    memory_target: 1342177280
    nls_language: 'AMERICAN'
    nls_territory: 'AMERICA'
    open_cursors: 300
    processes: 300
    remote_login_passwordfile: 'EXCLUSIVE'
  configTemplate: (unset)
  container: VCDO_1JL
  customEnvVars: (empty)
```

```

linked: false
logCollectionEnabled: false
mountBase: /mnt/provision
newDBID: false
nodeListeners: (empty)
operations:
  type: VirtualSourceOperations
  configureClone: (empty)
  postRefresh: (empty)
  postRollback: (empty)
  postSnapshot: (empty)
  postStart: (empty)
  postStop: (empty)
  preRefresh: (empty)
  preRollback: (empty)
  preSnapshot: (empty)
  preStart: (empty)
  preStop: (empty)
parentTdeKeystorePassword: *****
parentTdeKeystorePath: /u01/app/oracle/keystores/CDOMLOSR197/wallet
redoLogGroups: 3
redoLogSizeInMB: 50
reference: ORACLE_VIRTUAL_PDB_SOURCE-2
runtime:
  type: OraclePDBSourceRuntime
  accessible: true
  accessibleTimestamp: 2021-10-06T22:02:15.718Z
  activeInstances:
    0:
      type: OracleActiveInstance
      hostName: ip-10-110-234-67.delphix.com
      instanceName: CDOMLOSR197
      instanceNumber: 1
      databaseMode: READ_WRITE
      databaseRole: PRIMARY
      databaseSize: 913.4MB
      databaseStats: [ ... ]
      enabled: ENABLED
      lastNonLoggedLocation: 0
      status: RUNNING
runtimeMountInformation:
  type: UnixRuntimeMountInformation
  name: (unset)
  nfsVersion: 4
  nfsVersionReason: DEFAULT
staging: false
status: DEFAULT
tdeExportedKeyFileSecret: *****
tdeKeyIdentifier: AbSP7gninU+Gv1YQ/iEcJAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
tdeUUID: a3f26971-1df6-4c81-994f-4b2c582ded87
virtual: true

```

Operations

```

update
enable
disable
start
stop
upgrade

```

- Note that tdeKeyIdentifier is one of the last fields listed above. If we query the vPDB via sqlplus, we can see the matching key_id.

Note that any key generated by Delphix will include a tag with the format `dlpx_key_<tdeUUID>` .

```

SQL> alter session set container=VCDO_1JL;
Session altered.
SQL> select key_id, tag from v$encryption_keys;
KEY_ID
-----
TAG
-----
-
AbSP7gninU+Gv1YQ/iEcJAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
dlpx_key_a3f26971-1df6-4c81-994f-4b2c582ded87

```

- To generate a new unique encryption key, unset the value of tdeKeyIdentifier before a refresh or rewind operation.

```

delphix source 'VCDO_1JL'> update
delphix source 'VCDO_1JL' update *> unset tdeKeyIdentifier
delphix source 'VCDO_1JL' update *> ls
Properties
  type: OracleVirtualPdbSource
  name: VCDO_1JL
  allowAutoVDBRestartOnHostReboot: false
  customEnvVars: (empty)
  logCollectionEnabled: false
  newDBID: false
operations:
  type: VirtualSourceOperations
  configureClone: (empty)
  postRefresh: (empty)
  postRollback: (empty)
  postSnapshot: (empty)
  postStart: (empty)
  postStop: (empty)
  preRefresh: (empty)
  preRollback: (empty)
  preSnapshot: (empty)
  preStart: (empty)
  preStop: (empty)
parentTdeKeystorePassword: *****
parentTdeKeystorePath: /u01/app/oracle/keystores/CDOML0SR197/wallet

```

```

tdeKeyIdentifier: (unset) (*)
delphix source 'VCDO_1JL' update * > commit
Dispatched job JOB-18
SOURCE_UPDATE job started for "VCDO_1JL".
SOURCE_UPDATE job for "VCDO_1JL" completed successfully.

```

6. After the refresh or rewind, the new key identifier is now associated with vPDB that can be used for all future Delphix operations. View all the settings using the "ls" command.

```

delphix source 'VCDO_1JL' > ls
Properties
  type: OracleVirtualPdbSource
  name: VCDO_1JL
  allowAutoVDBRestartOnHostReboot: false
  archiveLogMode: true
  config: VCDO_1JL
  configParams:
    _cdb_disable_pdb_limit: TRUE
    audit_file_dest: '/u01/app/oracle/admin/CDOMLOSR197/adump'
    audit_trail: 'DB'
    compatible: '19.0.0'
    core_dump_dest: '/u01/app/oracle/diag/rdbms/cdomlosr197/CDOMLOSR197/
cdump'
    diagnostic_dest: '/u01/app/oracle'
    dispatchers: '(PROTOCOL=TCP) (SERVICE=CDOMLOSRA1DXDB)'
    enable_pluggable_database: TRUE
    log_archive_format: '%t_%s_%r.dbf'
    max_pdb: 4098
    memory_max_target: 1342177280
    memory_target: 1342177280
    nls_language: 'AMERICAN'
    nls_territory: 'AMERICA'
    open_cursors: 300
    processes: 300
    remote_login_passwordfile: 'EXCLUSIVE'
  configTemplate: (unset)
  container: VCDO_1JL
  customEnvVars: (empty)
  linked: false
  logCollectionEnabled: false
  mountBase: /mnt/provision
  newDBID: false
  nodeListeners: (empty)
  operations:
    type: VirtualSourceOperations
    configureClone: (empty)
    postRefresh: (empty)
    postRollback: (empty)
    postSnapshot: (empty)
    postStart: (empty)
    postStop: (empty)
    preRefresh: (empty)

```

```

preRollback: (empty)
preSnapshot: (empty)
preStart: (empty)
preStop: (empty)
parentTdeKeystorePassword: *****
parentTdeKeystorePath: /u01/app/oracle/keystores/CDOMLOS197/wallet
redoLogGroups: 3
redoLogSizeInMB: 50
reference: ORACLE_VIRTUAL_PDB_SOURCE-2
runtime:
  type: OraclePDBSourceRuntime
  accessible: true
  accessibleTimestamp: 2021-10-06T22:17:15.907Z
  activeInstances:
    0:
      type: OracleActiveInstance
      hostName: ip-10-110-234-67.delphix.com
      instanceName: CDOMLOS197
      instanceNumber: 1
      databaseMode: READ_WRITE
      databaseRole: PRIMARY
      databaseSize: 913.4MB
      databaseStats: [ ... ]
      enabled: ENABLED
      lastNonLoggedLocation: 0
      status: RUNNING
  runtimeMountInformation:
    type: UnixRuntimeMountInformation
    name: (unset)
    nfsVersion: 4
    nfsVersionReason: DEFAULT
  staging: false
  status: DEFAULT
  tdeExportedKeyFileSecret: *****
  tdeKeyIdentifier: AVEhXrBvmU+Cv+lK6ghT6oMAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
  tdeUUID: a3f26971-1df6-4c81-994f-4b2c582ded87
  virtual: true

```

- To specify a user-defined encryption key to be used for future Delphix operations, set `tdeKeyIdentifier` to the value of a valid `key_id` in the CDB's keystore.
Note that if an invalid `key_id` is provided, refresh or rewind will fail and it will be necessary to unset or update the `tdeKeyIdentifier` parameter with a valid `key_id`.
Note that this `key_id` will not have a corresponding `dlpx` tag unless it is a key previously generated by Delphix.

```

delphix source 'VCDO_1JL'> update
delphix source 'VCDO_1JL' update *> set tdeKeyIdentifier="AbSP7gninU+Gv1YQ/
iEcJAAAAAAAAAAAAAAAAAAAAAAAAAAAA"
delphix source 'VCDO_1JL' update *> ls
Properties
  type: OracleVirtualPdbSource
  name: VCDO_1JL

```

```

allowAutoVDBRestartOnHostReboot: false
customEnvVars: (empty)
logCollectionEnabled: false
newDBID: false
operations:
  type: VirtualSourceOperations
  configureClone: (empty)
  postRefresh: (empty)
  postRollback: (empty)
  postSnapshot: (empty)
  postStart: (empty)
  postStop: (empty)
  preRefresh: (empty)
  preRollback: (empty)
  preSnapshot: (empty)
  preStart: (empty)
  preStop: (empty)
parentTdeKeystorePassword: *****
parentTdeKeystorePath: /u01/app/oracle/keystores/CDOMLOS197/wallet
tdeKeyIdentifier: AbSP7gninU+Gv1YQ/iEcJAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
delphix source 'VCDO_1JL' update >> commit
Dispatched job JOB-22
SOURCE_UPDATE job started for "VCDO_1JL".
SOURCE_UPDATE job for "VCDO_1JL" completed successfully.

```

- After a refresh or rewind, this key identifier will be associated with vPDB and will be used for all future Delphix operations. View all the settings using the "ls" command.

```

delphix source 'VCDO_1JL'> ls
Properties
  type: OracleVirtualPdbSource
  name: VCDO_1JL
  allowAutoVDBRestartOnHostReboot: false
  archiveLogMode: true
  config: VCDO_1JL
  configParams:
    _cdb_disable_pdb_limit: TRUE
    audit_file_dest: '/u01/app/oracle/admin/CDOMLOS197/adump'
    audit_trail: 'DB'
    compatible: '19.0.0'
    core_dump_dest: '/u01/app/oracle/diag/rdbms/cdomlosr197/CDOMLOS197/
cdump'
    diagnostic_dest: '/u01/app/oracle'
    dispatchers: '(PROTOCOL=TCP) (SERVICE=CDOMLOS197CA1DXDB)'
    enable_pluggable_database: TRUE
    log_archive_format: '%t_%s_%r.dbf'
    max_pdbs: 4098
    memory_max_target: 1342177280
    memory_target: 1342177280
    nls_language: 'AMERICAN'
    nls_territory: 'AMERICA'
    open_cursors: 300

```

```

    processes: 300
    remote_login_passwordfile: 'EXCLUSIVE'
configTemplate: (unset)
container: VCDO_1JL
customEnvVars: (empty)
linked: false
logCollectionEnabled: false
mountBase: /mnt/provision
newDBID: false
nodeListeners: (empty)
operations:
  type: VirtualSourceOperations
  configureClone: (empty)
  postRefresh: (empty)
  postRollback: (empty)
  postSnapshot: (empty)
  postStart: (empty)
  postStop: (empty)
  preRefresh: (empty)
  preRollback: (empty)
  preSnapshot: (empty)
  preStart: (empty)
  preStop: (empty)
parentTdeKeystorePassword: *****
parentTdeKeystorePath: /u01/app/oracle/keystores/CDOMLOSR197/wallet
redoLogGroups: 3
redoLogSizeInMB: 50
reference: ORACLE_VIRTUAL_PDB_SOURCE-2
runtime:
  type: OraclePDBSourceRuntime
  accessible: true
  accessibleTimestamp: 2021-10-06T22:17:15.907Z
  activeInstances:
    0:
      type: OracleActiveInstance
      hostName: ip-10-110-234-67.delphix.com
      instanceName: CDOMLOSR197
      instanceNumber: 1
      databaseMode: READ_WRITE
      databaseRole: PRIMARY
      databaseSize: 913.4MB
      databaseStats: [ ... ]
      enabled: ENABLED
      lastNonLoggedLocation: 0
      status: RUNNING
runtimeMountInformation:
  type: UnixRuntimeMountInformation
  name: (unset)
  nfsVersion: 4
  nfsVersionReason: DEFAULT
staging: false
status: DEFAULT

```

```
tdeExportedKeyFileSecret: *****  
tdeKeyIdentifier: AbSP7gninU+Gv1YQ/iEcJAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
tdeUUID: a3f26971-1df6-4c81-994f-4b2c582ded87  
virtual: true
```

CLI Cookbook: Replication

These topics describe how to use the command-line interface for replication tasks.

This section covers the following topics:

- [CLI Cookbook: Adding a Replication Spec](#)
- [CLI Cookbook: Deleting a Replication Spec](#)
- [CLI Cookbook: Failing Over a Namespace](#)
- [CLI Cookbook: Mapping Replication Specs to Objects](#)
- [CLI Cookbook: Triggering Immediate Execution of a Replication Spec](#)

CLI Cookbook: Adding a Replication Spec

This topic describes how to use the command-line interface to add a replication specification to the Delphix Engine.

Unlike the GUI, the CLI supports the ability to manage multiple replication specifications within a single system. This allows updates to be sent to multiple systems from a single point.

Prerequisites

- You should review the topic [Replication Overview](#) to understand which objects are copied as part of a backup or restore operation, as well as the dependencies between objects.

Procedure

- Switch to the replication spec context.

```
delphix> cd replication/spec
delphix replication spec> ls
Operations
create
```

- Create a new replication spec.

```
delphix replication spec> create
delphix replication spec create *> ls
Properties
  type: ReplicationSpec
  name: (unset)
  bandwidthLimit: (unset)
  enabled: (unset)
  encrypted: (unset)
  objects: (required)
  schedule: (unset)
  targetCredential:
    type: PasswordCredential
    password: (required)
  targetHost: (required)
  targetPrincipal: (required)
```

- Specify the target hostname, user, and credentials.

```
delphix replication spec create *> set targetHost=exampleHost.mycompany.com
delphix replication spec create *> set targetPrincipal=delphix_admin
delphix replication spec create *> set targetCredential.password=password
```

SDD Secure replication

To create a Selective Data Distribution (SDD) type spec set the objectSpecification.type

```
delphix replication spec create *> set
objectSpecification.type=ReplicationSecureList
```

This parameter defaults to ReplicationList which implies a regular replication spec.

Target Principal

The target principal must be a Delphix user on the target host who has domain privileges.

4. Specify the set of objects to replicate.

- a. To replicate all dSources and VDBs on the system, specify ``DOMAIN`` as the list of objects.

```
delphix replication spec create *> set objectSpecification.objects=`DOMAIN`
```

- b. To replicate a subset of Groups, VDBs, dSources, and Timeflows specify their names as a comma-separated list.

```
delphix replication spec create *> set objectSpecification.objects=Group1/
vdb1,Group2/vdb2
```

Name Completion

The CLI will provide possible completions for all objects in the system, but only groups, dSources and VDBs can be specified. Attempts to replicate other types of objects will generate an error when the operation is committed.

SDD Secure replication

If `objectSpecification.type=ReplicationSecureList` was selected, then `objectSpecification.containers` needs to be used instead of `objectSpecification.objects`. The containers have to be Masked VDBs

5. Commit the operation.

```
delphix replication spec create *> commit
`REPLICATION_SPEC-1`
```

CLI Cookbook: Deleting a Replication Spec

This topic describes how to use the command-line interface to delete a replication spec.

Procedure

1. Switch to the replication spec context and list the specs on the system.

```
delphix> cd replication/spec
delphix replication spec> ls
Objects
REFERENCE          TARGETHOST
REPLICATION_SPEC-1 exampleHost.mycompany.com

Operations
create
```

2. Select the replication spec to remove.

```
delphix replication spec> select REPLICATION_SPEC-1
delphix replication spec "exampleHost.mycompany.com">
```

3. Remove the spec.

```
delphix replication spec "exampleHost.mycompany.com"> delete
delphix replication spec "exampleHost.mycompany.com" delete *> commit
```

CLI Cookbook: Failing Over a Namespace

This topic describes how to use the command line interface to fail over a namespace.

Procedure

1. Switch to the namespace context and list the namespaces on the system.

```
delphix> cd namespace
delphix namespace> ls
Objects
NAME
[172.16.203.93]

Operations
lookup
```

2. Select the namespace to failover.

```
delphix namespace> select [172.16.203.93]
delphix namespace "[172.16.203.93]">
```

3. Failover the namespace.

```
delphix namespace "[172.16.203.93]"> failover
delphix namespace "[172.16.203.93]" failover *> commit
```

Failover

Failover will sever the replication connection and make objects in the namespace part of the live system. This will prevent the target from receiving subsequent incremental updates.

CLI Cookbook: Mapping Replication Specs to Objects

After creating replication specs often you will want to see what is mapping to which target, this document will show you how to find that information in the CLI. It will also show you how to navigate the replication directory in the CLI.

1. ssh into the CLI as a user with delphix_admin privileges.

```
ssh admin@yourengine
```

2. Next go to the replication directory in the CLI and list all of the specs.

```
delphix > replication
delphix replication > spec
delphix replication spec > ls
Objects

REFERENCE          TARGETHOST

REPLICATION_SPEC-1 test1

Operations

create
```

3. Then select the replication spec to find where it maps to and what objects are selected with it.

```
delphix replication spec> select REPLICATION_SPEC-1

delphix replication spec 'test'> ls

Properties

  type: ReplicationSpec

  name: test

  bandwidthLimit: 0

  description: (unset)

  enabled: false

  encrypted: false

  numberOfConnections: 1
```

```
objectSpecification:
  type: ReplicationList
  name: (unset)
  objects: Untitled/dbdhcp3
reference: REPLICATION_SPEC-1
runtime:
  type: ReplicationSpecRuntime
schedule: (unset)
tag: 0ddae174-9486-4363-9704-bfc3398e547e
targetCredential:
  type: PasswordCredential
  password: *****
targetHost: test1
targetPort: 8415
targetPrincipal: delphix
useSystemSocksSetting: false
```

Operations

delete

update

execute

CLI Cookbook: Triggering Immediate Execution of a Replication Spec

This topic describes how to use the command-line interface to trigger an immediate execution of a replication spec in the Delphix Engine.

Procedure

1. Switch to the replication spec context and list the specs on the system.

```
delphix> cd replication/spec
delphix replication spec> ls
Objects
REFERENCE          TARGETHOST
REPLICATION_SPEC-1 exampleHost.mycompany.com

Operations
create
```

2. Select the replication spec to execute.

```
delphix replication spec> select REPLICATION_SPEC-1
delphix replication spec "exampleHost.mycompany.com">
```

3. Execute the spec.

```
delphix replication spec "exampleHost.mycompany.com"> execute
delphix replication spec "exampleHost.mycompany.com" execute *> commit
  Dispatched job JOB-7
  REPLICATION_SEND job started.
  Connecting to target "exampleHost.mycompany.com".
  Preparing replication update.
  Starting incremental replication update.
  Sending metadata.
  Sending data for "Untitled".
  Sending data for "Untitled/redsox1".
  Transfer completed in 0:00:01, sent 1.39MB (1.39MB/s).
  Committing serialization state.
  REPLICATION_SEND job completed successfully.
```

CLI Cookbook: Delphix Self-Service Actions

These entries will help with some of the Delphix Self-Service Actions that can be performed in the CLI and therefore with the Delphix APIs:

- [CLI Cookbook: How to Create a Delphix Self-Service Bookmark](#)
- [CLI Cookbook: How to Create a Delphix Self-Service Branch](#)
- [CLI Cookbook: How to Create a Delphix Self-Service Container](#)
- [CLI Cookbook: How to Create a Delphix Self-Service Database Template](#)
- [CLI Cookbook: How to Create a Delphix Self-Service User](#)
- [CLI Cookbook: How to Delete a Delphix Self-Service Bookmark](#)
- [CLI Cookbook: How to Delete a Delphix Self-Service Container](#)
- [CLI Cookbook: How to Delete a Delphix Self-Service Template](#)
- [CLI Cookbook: How to Refresh a Delphix Self-Service Container](#)
- [CLI Cookbook: How to Share a Delphix Self-Service Bookmark](#)
- [CLI Cookbook: How to Update a Delphix Self-Service Bookmark](#)

CLI Cookbook: How to Create a Delphix Self-Service Bookmark

Prerequisites:

- Have Delphix Self-Service user privileges.
- Know the branch you would like to create the bookmark on.
- Know the container or template that the branch belongs to.
- (Optional) Know when you would like the bookmark to expire.

Delphix Self-Service administrators can use this CLI cookbook recipe to create a bookmark on Delphix Self-Service using the Delphix Engine CLI.

 The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createBookmark.sh](#).

Creating a bookmark in Delphix Self-Service

```
#!/bin/bash
# A sample script for calls to the CLI. This one creates a Jet Stream bookmark.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
+Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#

##example##
./createBookmark.sh -e "2016-07-27T23:38:56.453Z" -t "2016-07-27T01:45:56.453Z" -T
[tag1,tag2,tag3,tag4,tag5] bkmrk3 JS_BRANCH-41

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="ars-6010.dlpxdc.co"
username="admin"

##### Functions
```

```

# Help Menu
function usage {
    echo "Usage: createBookmark.sh [[-h] | options...] <bookmark name> <branch name
format JS_BRANCH-n>"
    echo "Create a Jet Stream Bookmark on the given branch."
    echo ""
    echo "Positional arguments"
    echo " <bookmark name> "
    echo " <branch name>"
    echo ""
    echo "Optional Arguments:"
    echo " -h                Show this message and exit"
    echo " -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo " -u                Server user. Password needs to manually provide at run
time, otherwise revert to default"
    echo " -D                Description of this bookmark. Type: string"
    echo " -s                Pass [-s true] if need to make bookmark in shared
mode"
    echo " -e                The time at which the bookmark should be expired"
    echo " -t                The time at which the bookmark should be created. If no
time is included, the bookmark will be created at the latest point in time. Type:
date, must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
    echo " -T                A set of user-defined labels for this bookmark. No
spaces allowed. Array of Type: string. In format, [tag1,tag2,..] "
}

# Create Our Session, including establishing the API version.
function create_session
{
    echo "creating session..."
    SSH_CMD="ssh ${username}@${engine}"
    ${SSH_CMD} "version $VERSION"
    check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        echo $result
        exit 1
    fi
}

function get_branch

```

```

{
  echo "retrieving branch $branchRef to find Source Data Layout..."
  result=$(SSH_CMD "selfservice branch; ls; select ${branchRef}; ls")

  # Get everything in the result that comes after dataLayout.
  temp=${result#*dataLayout: }

  # Get rid of everything after
  dataLayout=${temp%% *}
}

function create_bookmark
{
  get_branch

  echo "creating bookmark..."

  # If there is not creation time, we need to use JSTimelinePointLatestTimeInput.
  if [ -z $creationTime ]
  then
    pointParams="set timelinePointParameters.sourceDataLayout=$dataLayout;"
    pointParams="$pointParams set
timelinePointParameters.type=JSTimelinePointLatestTimeInput;"
  else
    pointParams="set timelinePointParameters.sourceDataLayout=$dataLayout;"
    pointParams="$pointParams set
timelinePointParameters.type=JSTimelinePointTimeInput;"
    pointParams="$pointParams set
timelinePointParameters.time=\"$creationTime\";"
    pointParams="$pointParams set timelinePointParameters.branch=\"$branchRef\";"

  fi

  if [[ -n $expirationTime ]]
  then
    pointParams="$pointParams set bookmark.expiration=\"$expirationTime\";"
  fi

  # These are the required parameters.
  paramString="
  \"bookmark\": {
    \"branch\": \"${branchRef}\",
    \"name\": \"${bookmarkName}\",
    paramString="selfservice bookmark; create; set bookmark.name=$bookmarkName; set
bookmark.branch=$branchRef;"
    paramString="$paramString $pointParams"

  # Fill in optional parameters if there are any.
  if [[ -n $description ]]

```

```

then
    paramString="$paramString set bookmark.description=\"$description\";"
fi

if [[ -n $shared ]]
then
    paramString="$paramString set bookmark.shared=$shared;"
fi

if [[ -n $tags ]]
then
    # Add quotes back to the passed in tags so they are processed correctly.
    tags=${tags//[/[\\]}
    tags=${tags//,/,\""}
    tags=${tags//]/\""}

    paramString="$paramString set bookmark.tags=$tags;"
fi
paramString="$paramString commit;"
result=$(SSH_CMD $paramString)
check_result
echo "Verifying job status..."
# Get everything in the result that comes after job.
temp=${result#*job}
# Get rid of everything after
resultArray=( $temp )
jobRef=( $resultArray )
jobString="job;select $jobRef;ls"
result=$(SSH_CMD $jobString)
check_result
# Get everything in the result that comes after job.
temp=${result#*jobState:}
# Get rid of everything after
resultArray=( $temp )
jobState=( $resultArray )

if [ $jobState = "COMPLETED" ]
then
    echo "successfully created bookmark $bookmarkName"
else
    echo "unable to create bookmark"
    echo result
fi
}

##### Main

while getopts "u:d:D:s:e:t:T:h" flag; do
    case "$flag" in
        u )                username=${OPTARG%:*}

```

```

    d )      ;;
             engine=$OPTARG
    D )      ;;
             description=$OPTARG
    s )      ;;
             shared=true
    t )      ;;
             creationTime=$OPTARG
    e )      ;;
             expirationTime=$OPTARG
    T )      ;;
             tags=$OPTARG
    h )      ;;
             usage
             exit
    * )      ;;
             usage
             exit 1
esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 2 positional arguments
if [ $# != 2 ]
then
    usage
    exit 1
fi

# Get the two positional arguments
bookmarkName=$1
shift
branchRef=$1
create_session
create_bookmark

```

CLI Cookbook: How to Create a Delphix Self-Service Branch

Delphix Self-Service administrators can use this CLI cookbook recipe to create a branch on Delphix Self-Service using the Delphix Engine CLI.

i The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createBranch.sh](#).

Creating a branch in Delphix Self-Service

```
#!/bin/bash

# A sample script for calls to the CLI. This one creates a Jet Stream branch.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#
# Note that the CLI only allows branches to be created from below
# 1) From latest point in time
# 2) From specific bookmark
# 3) From specific point in time

##examples##
# Create branch from latest point in time
#./createBranch.sh NewBranchName JS_DATA_CONTAINER-20
# Create branch from specific bookmark
#./createBranch.sh -b ExistingBookmarkName NewBranchName JS_DATA_CONTAINER
# Create branch from specific point in time
#./createBranch.sh -t "2016-07-27T01:45:56.453Z" -B JS_BRANCH-2 NewBranchName
JS_DATA_CONTAINER-20

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="ars3-6010.dlpxdc.co"
username="admin"
```

```
##### Functions

# Help Menu
function usage {
    echo "Usage: createBranch.sh [[-h] | options...] <name> <container>"
    echo "Create a Jet Stream Bookmark on the given branch."
    echo ""
    echo "Positional arguments"
    echo "  <NewBranchName>"
    echo "  <container> format JS_DATA_CONTAINER-<n>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h                Show this message and exit"
    echo "  -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo "  -u                Server user. Password needs to manually provide at run
time, otherwise revert to default"
    echo "  -b                Bookmark name from which need to create branch. If no
bookmark is included, the branch will be created at the latest point in time. Type:
string"
    echo "  -t                The time at which the branch should be created. If no
time is included, the branch will be created at the latest point in time. Type: date,
must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
}

# Create Our Session, including establishing the API version.
function create_session
{
    echo "creating session..."
    SSH_CMD="ssh ${username}@${engine}"
    ${SSH_CMD} "version $VERSION"
    check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        echo $result
        exit 1
    fi
}

```

```

function create_branch
{

    # If there is not timeInput and no bookmark name, we need to use
    JSTimelinePointLatestTimeInput.
    if [ -z $inputTime ] && [ -z $bookmark ]
    then
        #code to use JSTimelinePointLatestTimeInput
        pointParams="edit timelinePointParameters; set
type=JSTimelinePointLatestTimeInput;"
        pointParams="$pointParams set sourceDataLayout=$container; commit;"
    # If there is a timeInput and no bookmark name, we need to use Input Time.
    elif [ -n $inputTime ] && [ -z $bookmark ]
    then
        #code to use JSTimelinePointTimeInput
        pointParams="edit timelinePointParameters; set
type=JSTimelinePointTimeInput;"
        pointParams="$pointParams set time=$inputTime; set branch=$branch; commit;"
    # If there is a bookmark name and no time input, we need to use bookmark
    elif [ -z $inputTime ] && [ -n $bookmark ]
    then
        #code to use JSTimelinePointBookmarkInput
        pointParams="set timelinePointParameters.bookmark=$bookmark;"
        pointParams="$pointParams commit;"
    fi

    # These are the required parameters.
    paramString="selfservice branch create;set name=$branchName; set
dataContainer=$container;"

    #Add additional optional parameter
    paramString="$paramString $pointParams"
    #echo $paramString
    echo "Creating Branch..."
    result=$((${SSH_CMD} $paramString)
    check_result

    echo "Verifying job status..."
    # Get everything in the result that comes after job.
        temp=${result#*job}
    # Get rid of everything after
    resultArray=( $temp)
    jobRef=( $resultArray)
        jobString="job;select $jobRef;ls"
        result=$((${SSH_CMD} $jobString)
    check_result
    # Get everything in the result that comes after job.
    temp=${result#*jobState:}
    # Get rid of everything after
    resultArray=( $temp)
    jobState=( $resultArray)
        if [ $jobState = "COMPLETED" ]

```

```

    then
        echo "Successfully created branch $branchName"
    else
        echo "Unable to create branch"
        echo $result
    fi
}

##### Main

while getopts "u:d:b:t:B:h" flag; do
    case "$flag" in
        u )          username=${OPTARG%:*}
                    ;;
        d )          engine=$OPTARG
                    ;;
        b )          bookmark=$OPTARG
                    ;;
        B )          branch=$OPTARG
                    ;;
        t )          inputTime=$OPTARG
                    ;;
        h )          usage
                    exit
                    ;;
        * )          usage
                    exit 1
    esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 2 positional arguments
if [ $# != 2 ]
then
    usage
    exit 1
fi

# Get the two positional arguments
branchName=$1
shift
container=$1

create_session
create_branch

```

CLI Cookbook: How to Create a Delphix Self-Service Container

Delphix Self-Service administrators can use this CLI cookbook recipe to create a container on Delphix Self-Service using the Delphix Engine CLI.

 The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createContainer.sh](#).

Creating a container in Delphix Self-Service

```
#!/bin/bash

# A sample script for calls to the CLI. This one creates a Jet Stream container.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="name-of-engine.dlpxdc.co"
username="admin"

##examples##
# Create container from latest point in time
#./createContainer.sh -n "testsource" testcont ORACLE_DB_CONTAINER-269
JS_DATA_TEMPLATE-13
# Create container from specific bookmark
#./createContainer.sh -n "testsource" -b JS_BOOKMARK-77 testcont
ORACLE_DB_CONTAINER-269 JS_DATA_TEMPLATE-13
# Create container from specific point in time
#./createContainer.sh -n "testsource" -t "2016-08-08T10:00:00.000Z" -B JS_BRANCH-50
testcont ORACLE_DB_CONTAINER-269 JS_DATA_TEMPLATE-13

#NOTE: this script will add one container and assign one owner for the container.
```

```

##### Functions

# Help Menu
function usage {
    echo "Usage: createContainer.sh [[-h] | options...] <containername> <vdb>
<template>"
    echo "Create a Jet Stream Container."
    echo ""
    echo "Positional arguments"
    echo "  <name>"
    echo "  <container> format JS_DATA_CONTAINER-<n>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h                Show this message and exit"
    echo "  -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo "  -u                Server user. Password needs to manually provide at run
time, otherwise revert to default"
    echo "  -n                SourceName need to display for container.
(Mandatory)"
    echo "  -b                Bookmark name from which need to create container. If
no bookmark is included, the branch will be created at the latest point in time.
Type: string. Format JS_BOOKMARK-<n> (Optional)"
    echo "  -B                Branch reference from which we need to pick up time
from where the container should be created. Type: string. Format JS_BRANCH-<n>"
    echo "  -t                The time at which the branch should be created. This
must be accompanied with branch name from which need to pick up time. Type: date,
must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
    echo "  -N                Optional container notes, if need to add any. Type:
String"
    echo "  -o                Optional owner, to whom we need to assign this
container. Type: String. Format USER-<n>"
}

# Create Our Session, including establishing the API version.
function create_session
{
    echo "creating session..."
    SSH_CMD="ssh ${username}@${engine}"
    ${SSH_CMD} "version $VERSION"
    check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        echo $result
        exit 1
    fi
}

```

```

    fi
}

function create_container
{
    # If there is not timeInput and no bookmark name, we need to use
    JSTimelinePointLatestTimeInput.
    if [[ -z $inputTime && -z $bookmark ]]
    then
        pointParams="set
timelinePointParameters.type=JSTimelinePointLatestTimeInput;"
        pointParams="$pointParams set
timelinePointParameters.sourceDataLayout=$template;"

        # If there is a timeInput, no bookmark name and a branch name, we need to use
        Input Time.
        elif [[ -n $inputTime && -z $bookmark && -n $branch ]]
        then
            pointParams="set timelinePointParameters.type=JSTimelinePointTimeInput;"
            pointParams="$pointParams set timelinePointParameters.time=\"${inputTime}\";"
            pointParams="$pointParams set timelinePointParameters.branch=$branch;"
        # If there is a bookmark name and no time input, we need to use bookmark
        elif [[ -z $inputTime && -n $bookmark ]]
        then
            pointParams="set timelinePointParameters.type=JSTimelinePointBookmarkInput;"
            pointParams="$pointParams set timelinePointParameters.bookmark=\"${bookmark}
\";"
        else
            usage
            exit 1
        fi

        # These are the required parameters.

        paramString="selfservice container create;set name=\"${containerName}\";"
        paramString="$paramString set template=\"${template}\";"

        if [[ -n $containerNotes ]]
        then
            paramString="$paramString set notes=\"${containerNotes}\";"
        fi

        if [[ -n $owners ]]
        then
            paramString="$paramString set owner=\"${owners}\";"
        fi
        paramString="$paramString $pointParams;"
        paramString="$paramString edit dataSources; add; set container=$VDB;"
        paramString="$paramString edit source; set type=JSDataSource; set priority=1;set
name=\"${sourceName}\";"

```

```

if [[ -n $sourcedesc ]]
then
    paramString="$paramString set description=\"${sourcedesc}\";"
fi

paramString="$paramString commit;"
#echo $paramString
echo "Creating Container..."
result=$((${SSH_CMD} $paramString)
check_result

echo "Verifying job status..."
# Get everything in the result that comes after job.
temp=${result#*job}
# Get rid of everything after
resultArray=( $temp)
jobRef=( $resultArray)
jobString="job;select $jobRef;ls"
result=$((${SSH_CMD} $jobString)
check_result
# Get everything in the result that comes after job.
temp=${result#*jobState:}
# Get rid of everything after
resultArray=( $temp)
jobState=( $resultArray)
    if [ $jobState = "COMPLETED" ]
then
    echo "Successfully created Container $containerName"
else
    echo "Unable to create Container"
    echo $result
fi
}

##### Main

while getopts "u:d:b:t:B:D:n:N:o:h" flag; do
    case "$flag" in
        u )           username=${OPTARG%:*}
                    ;;
        d )           engine=$OPTARG
                    ;;
        b )           bookmark=$OPTARG
                    ;;
        t )           inputTime=$OPTARG
                    ;;
        D )           sourcedesc=$OPTARG
                    ;;
        n )           sourceName=$OPTARG
                    ;;
        B )           branch=$OPTARG
    
```

```
        N ) ;;
            containerNotes=$OPTARG
            ;;
        o )
            owners=$OPTARG
            ;;
        h )
            usage
            exit
            ;;
        * )
            usage
            exit 1
    esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 3 positional arguments
if [ $# != 3 ]
then
    usage
    exit 1
fi

# Get the three positional arguments
containerName=$1
shift
VDB=$1
shift
template=$1

create_session
create_container
```

CLI Cookbook: How to Create a Delphix Self-Service Database Template

Delphix Self-Service administrators can use this CLI cookbook recipe to create a database template on Delphix Self-Service using the Delphix Engine CLI.

 The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createDBTemplate.sh](#).

Creating a database template in Delphix Self-Service

```
#!/bin/bash

# A sample script for calls to the CLI. This one creates a Jet Stream Template.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#
#

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="ars-6010.dlpxdc.co"
username="admin"

##examples##
# Create template with mandatory params
# Create template with adding optional params, Notes and Description
#./createDBTemplate.sh -n <sourceName> -N "<templateNotes>" -D "<AnyDescription>"
<templateName> <containerName>

## NOTE: This script is to add one source per template and it will not add any
properties for template, container or source.

##### Functions

# Help Menu
function usage {
```

```

    echo "Usage: createDBTemplate.sh [[-h] | options...] <template_name>
<source_container>"
    echo "Create a Jet Stream Dat Template."
    echo ""
    echo "Positional arguments"
    echo "  <template_name>"
    echo "  <source_container>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h                Show this message and exit"
    echo "  -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo "  -u                Server user. Password needs to manually provide at run time,
otherwise revert to default"
    echo "  -n                source name to display on JS template screen"
    echo "  -N                template notes, if any. Type: String"
    echo "  -D                source description, if any. Type: String"
}

# Create Our Session, including establishing the API version.
function create_session
{
    echo "creating session..."
    SSH_CMD="ssh ${username}@${engine}"
    ${SSH_CMD} "version $VERSION"
    check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        echo $result
        exit 1
    fi
}

function create_template
{
    paramString="selfservice template create;set name=$templateName;"
    if [[ -n $templatenotes ]]
    then
        paramString="$paramString set notes=\"$templatenotes\";"
    fi
    paramString="$paramString edit dataSources;add;"
    paramString="$paramString set container=$sourceContainer;set
source.name=$sourcename;set source.priority=1;"
}

```

```

if [[ -n $sourcedesc ]]
then
    paramString="$paramString set source.description=\"$sourcedesc\";"
fi
paramString="$paramString commit;"
#echo $paramString
echo "Creating Data Template..."
result=$(SSH_CMD $paramString)
check_result
echo "New JetStream template $templateName successfully created"
}

##### Main

while getopt "u:d:n:N:D:h" flag; do
    case "$flag" in
        u )           username=${OPTARG%:*}
                    ;;
        d )           engine=$OPTARG
                    ;;
        n )           sourcename=$OPTARG
                    ;;
        N )           templatenotes=$OPTARG
                    ;;
        D )           sourcedesc=$OPTARG
                    ;;
        h )           usage
                    exit 1
                    ;;
        * )           usage
                    exit 1
    esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 2 positional arguments
if [ $# != 2 ]
then
    usage
    exit 1
fi

# Get the two positional arguments
templateName=$1
shift
sourceContainer=$1

```

```
create_session  
create_template
```

CLI Cookbook: How to Create a Delphix Self-Service User

Delphix Self-Service administrators can use this CLI cookbook recipe to create a user on Delphix Self-Service using the Delphix Engine CLI.

i The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createJSUser.sh](#).

Creating a Delphix Self-Service User

```
#!/bin/bash
# A sample script for calls to the CLI. This one creates a Self-Service user.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS533/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#
# Note that the CLI only allows branches to be created from existing bookmarks.
##### Constants
# Describes a Delphix software revision.
VERSION="1.6.2"

##### Default Values. These can be overwritten with optional arguments.
engine="10.0.1.10"
username="dev"

##examples##
# Create user with NATIVE authentication
#./createJSUser.sh -P <password> NATIVE <username>
# Create user with LDAP authentication
#./createJSUser.sh -r <principal> <LDAP username>

##### Functions
# Help Menu
function usage {
echo "Usage: createJSUser.sh [[-h] | options...] <auth> <newjsuser>"
echo "Create a Self-Service Only user."
echo ""
echo "Positional arguments"
echo " <auth type NATIVE/LDAP>"
echo " <newjsuser>"
echo ""
echo "Optional Arguments:"
echo " -h Show this message and exit"
echo " -d Delphix engine IP address or host name, otherwise revert to default"
```

```

echo " -u Server user. Password needs to manually provide at run time, otherwise
revert to default"
echo " -P password for NATIVE authentication, MUST incase auth=NATIVE"
echo " -f firstName of user"
echo " -l lastName of user"
echo " -e emailAddress of user"
echo " -o homePhoneNumber of user"
echo " -m mobilePhoneNumber of user"
echo " -w workPhoneNumber of user"
echo " -r principal for LDAP authentication, MUST incase of auth=LDAP"
}

# Create Our Session, including establishing the API version.
function create_session
{
echo "creating session..."
SSH_CMD="ssh ${username}@${engine}"
${SSH_CMD} "version $VERSION"
check_result
}
# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
exitStatus=$?
if [ $exitStatus -ne 0 ]
then
echo "command failed with exit status $exitStatus"
echo $result
exit 1
fi
}

function create_user
{
# Check on authorization type
paramString="user create;"
if [[ $authtype = "NATIVE" && -n $userpwd ]]
then
pointParams="set authenticationType=$authtype;"
pointParams="$pointParams set credential.type=PasswordCredential; set
credential.password=$userpwd;"
elif [[ $authtype = "LDAP" && -n $principal ]]
then
pointParams="set authenticationType=$authtype; set principal=$principal;"
fi
# These are the required parameters.
paramString="$paramString set type=User; set name=$newjsuser;"

# Fill in optional parameters if there are any.
if [[ -n $firstname ]]
then
paramString="$paramString set firstName=\"$firstname\";"

```

```

fi

if [[ -n $lastname ]]
then
paramString="$paramString set lastName=\"$lastname\";"
fi

if [[ -n $emailaddress ]]
then
paramString="$paramString set emailAddress=\"$emailaddress\";"
fi

if [[ -n $homephone ]]
then
paramString="$paramString set homePhoneNumber=\"$homephone\";"
fi

if [[ -n $mobilephone ]]
then
paramString="$paramString set mobilePhoneNumber=\"$mobilephone\";"
fi

if [[ -n $workphone ]]
then
paramString="$paramString set workPhoneNumber=\"$workphone\";"
fi
paramString="$paramString ${pointParams} commit;"
#echo $paramString
echo "Creating user..."
result=$((${SSH_CMD} $paramString)
check_result
#echo $result
echo "New user $newjsuser successfully created"

##### ROLE-3 is Self-Service Role
paramString="authorization create;"
paramString="$paramString set type=Authorization; set role=ROLE-3; set
target=$newjsuser; set user=$newjsuser;commit;"
#echo $paramString
result=$((${SSH_CMD} $paramString)
check_result
echo "Assigned Self-Service Role to user $newjsuser"
}

##### Main
##### Main
while getopts "u:d:P:r:f:l:e:o:m:w:h" flag; do
case "$flag" in
u ) username=${OPTARG%:*}
;;
d ) engine=$OPTARG
;;

```

```
P ) userpwd=$OPTARG
;;
r ) principal=$OPTARG
;;
f ) firstname=$OPTARG
;;
l ) lastname=$OPTARG
;;
e ) emailaddress=$OPTARG
;;
o ) homephone=$OPTARG
;;
m ) mobilephone=$OPTARG
;;
w ) workphone=$OPTARG
;;
h ) usage
exit
;;
* ) usage
exit 1

esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))
# Check that there are 2 positional arguments
if [ $# != 2 ]
then
echo "usage1"
usage
exit 1
fi
# Get the two positional arguments
authtype=$1
shift
newjsuser=$1
create_session
create_user
```

CLI Cookbook: How to Delete a Delphix Self-Service Bookmark

Prerequisites:

- Have Delphix Self-Service user privileges
- Know the bookmark you would like to delete

Procedure:

1. Log into Delphix Engine as a Delphix Self-Service user or admin.

```
ssh jsUser@<yourengine>
```

2. Navigate to the Delphix Self-Service bookmarks, and choose the one you would like to delete.

```
jsUser > selfservice  
bookmark  
jsUser selfservice bookmark > ls  
jsUser selfservice bookmark > select <jsBookmark>
```

3. Delete the bookmark.

```
jsUser selfservice bookmark <jsBookmark> *> delete  
jsUser selfservice bookmark <jsBookmark> delete *> commit
```

CLI Cookbook: How to Delete a Delphix Self-Service Container

Prerequisites:

- Know the container you want to delete
- Have delphix_admin privileges

Procedure:

1. Log into Delphix Engine as admin

```
ssh delphix@<yourengine>
```

2. Navigate to the Delphix Self-Service container that you want to delete

```
delphix > selfservice container  
delphix selfservice container > ls  
delphix selfservice container > select CONTAINER_X  
delphix selfservice container 'CONTAINER_X' >
```

3. Delete container, **note you need to set if you want to delete the VDBs in the container (false will preserve the VDBs and true the VDBs will be deleted along with the container)**

```
delphix selfservice container 'CONTAINER_X' > delete  
delphix selfservice container 'CONTAINER_X' delete * > set deleteDataSources=<true/false>  
delphix selfservice container 'CONTAINER_X' delete * > commit
```

CLI Cookbook: How to Delete a Delphix Self-Service Template

Prerequisites:

- Template has no dependant containers
- Know the name of the template you are going to delete
- Have delphix_admin privileges (note JetStream Only users and Delphix GUI Owners cannot delete templates)

Procedure:

1. Log into your Delphix Engine using delphix_admin (or admin privileged account)

```
ssh delphix_admin@<youengine>
```

2. Find the template you want to delete

```
delphix > selfservice  
delphix selfservice template > ls  
delphix selfservice template > select 'TEMPLATE_X'  
delphix selfservice template 'TEMPLATE_X'>
```

3. Delete template and commit

```
delphix selfservice template 'TEMPLATE_X'> delete  
delphix selfservice template 'TEMPLATE_X' delete *> commit
```

CLI Cookbook: How to Refresh a Delphix Self-Service Container

i The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [refreshContainer.sh](#).

Refreshing a container in Delphix Self-Service

```
#!/bin/bash

# A sample script for calls to the CLI. This one refresh Jet Stream container.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.4"

##### Default Values. These can be overwritten with optional arguments.
engine="de-6041-doc-938-vn.dlpxdc.co"
username="admin"

##examples##
# Refresh container from latest point in time of Template
#./refreshContainer.sh -T JS_DATA_TEMPLATE-13 JS_DATA_CONTAINER-20
# Refresh container from specific bookmark
#./refreshContainer.sh -b JS_BOOKMARK-76 JS_DATA_CONTAINER-20
# Refresh container from specific point in time of branch
#./refreshContainer.sh -t "2016-08-08T10:00:00.000Z" -B JS_BRANCH-4
JS_DATA_CONTAINER-20

##### Functions

# Help Menu
function usage {
echo "Usage: refreshContainer.sh [[-h] | options...] <containername>"
echo "Create a Jet Stream Bookmark on the given branch."
}
```

```

echo ""
echo " You need to specify either -T, -B, -b or -t to refresh container. With -t
option you also need to specify -B also"
echo "Positional arguments"
echo " <containerName>"
echo ""
echo "Optional Arguments:"
echo " -h Show this message and exit"
echo " -d Delphix engine IP address or host name, otherwise revert to default"
echo " -u Server user. Password needs to manually provide at run time, otherwise
revert to default"
echo " -T template reference from which need to refresh from latest point in time"
echo " -B Branch reference from which we need to pick up time from where the
container should be refreshed. Type: string. Format JS_BRANCH-<n> (Optional)"
echo " -b Bookmark name from which need to refresh container. If no bookmark is
included, the branch will be created at the latest point in time. Type: string.
Format JS_BOOKMARK-<n> (Optional)"
echo " -t The time from where the container should be refreshed. This must be
accompanied with branch reference from which need to pick up time. Type: date, must
be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
}

# Create Our Session, including establishing the API version.
function create_session
{
echo "Creating session ..."
SSH_CMD="ssh ${username}@${engine}"
${SSH_CMD} "version $VERSION"
check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
exitStatus=$?
if [ $exitStatus -ne 0 ]
then
echo "command failed with exit status $exitStatus"
echo $result
exit 1
fi
}

function restore_container
{
paramString="selfservice/container/select ${containerRef}; restore;"

# If there is no time input and no bookmark name, we need to use
JSTimelinePointLatestTimeInput from template.
if [[ -n $template && -z $inputTime && -z $bookmark ]]

```

```

then
pointParams="edit timelinePointParameters; set type=JSTimelinePointLatestTimeInput;
set sourceDataLayout=\"${template}\";"
# If there is a timeInput and no bookmark name, we need to use Input Time.
elif [[ -n $inputTime && -n $branch && -z $bookmark && -z $template ]]
then
pointParams="edit timelinePointParameters; set type=JSTimelinePointTimeInput; set
time=\"${inputTime}\"; set branch=\"${branch}\";"
# If there is a bookmark name and no time input, we need to use bookmark
elif [[ -n $bookmark && -z $template && -z $inputTime ]]
then
pointParams="edit timelinePointParameters; set type=JSTimelinePointBookmarkInput; set
bookmark=\"${bookmark}\";"
else
usage
exit 1
fi

paramString="$paramString $pointParams commit;"
echo "Executing the command on CLI ..."
echo $paramString
result=$((${SSH_CMD} $paramString)
check_result

echo "Verifying job status ..."
# Get everything in the result that comes after job.
temp=${result#*job}
# Get rid of everything after
resultArray=( $temp)
jobRef=( $resultArray)
jobString="job;select $jobRef;ls"
result=$((${SSH_CMD} $jobString)
check_result
# Get everything in the result that comes after job.
temp=${result#*jobState:}
# Get rid of everything after
resultArray=( $temp)
jobState=( $resultArray)

if [ $jobState = "COMPLETED" ]
then
echo "Successfully refreshed container- [$containerRef]."
else
echo "Unable to refresh container- [$containerRef]."
echo result
fi

}

##### Main

```

```
while getopts "u:d:T:b:t:B:h" flag; do
case "$flag" in
u ) username=${OPTARG%:*}
;;
d ) engine=$OPTARG
;;
T ) template=$OPTARG
;;
b ) bookmark=$OPTARG
;;
t ) inputTime=$OPTARG
;;
B ) branch=$OPTARG
;;
h ) usage
exit
;;
* ) usage
exit 1
esac

done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 1 positional arguments
if [ $# != 1 ]
then
usage
exit 1
fi

# Get the one positional arguments
containerRef=$1

create_session
restore_container
```

CLI Cookbook: How to Share a Delphix Self-Service Bookmark

Prerequisites:

- Have Delphix Self-Service user privileges.
- Know the bookmark you would like to share.

Delphix Self-Service administrators can use this CLI cookbook recipe to share a bookmark on Delphix Self-Service using the Delphix Engine CLI.

i The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [shareBookmark.sh](#).

Sharing a bookmark in Delphix Self-Service

```
#!/bin/bash

# A sample script for calls to the CLI. This one shares Bookmark across containers in
# same template.
#
# VERY IMPORTANT: In order for this to work, you need to go through the steps here:
# https://docs.delphix.com/display/DOCS43/CLI+Cookbook%3A+Configuring+Key-
# Based+SSH+Authentication+for+Automation
# After this you will not need to use a username and password to log into the delphix
# engine. If you do not
# setup the SSH authentication you will have to manually enter the password.
#

##examples##
# Share Bookmark
#./shareBookmark.sh -a share JS_BOOKMARK-75
# Unshare Bookmark
#./shareBookmark.sh -a unshare JS_BOOKMARK-75

##### Constants

# Describes a Delphix software revision.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="10.110.213.109"
```

```

username="admin"

##### Functions

# Help Menu
function usage {
    echo "Usage: shareBookmark.sh [[-h] | options...] -a share/unshare
<bookmarkName>"
    echo "Share/Unshare JetStream bookmark"
    echo ""
    echo "Positional arguments"
    echo "bookmarkName. Format: JS_BOOKMARK-<n>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h                Show this message and exit"
    echo "  -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo "  -u                Server user. Password needs to manually provide at run
time, otherwise revert to default"
    echo "  -a                action to perform on bookmark. Type:String.
Values:share/unshare"
}

# Create Our Session, including establishing the API version.
function create_session
{
    echo "creating session..."
    SSH_CMD="ssh ${username}@${engine}"
    ${SSH_CMD} "version $VERSION"
    check_result
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        echo $result
        exit 1
    fi
}

function bookmark_action
{
    # Change share mode of bookmark

```

```

echo "Changing share mode.."
if [[ $action = "share" ]]
then
    paramString="selfservice bookmark select $bookmarkName share; commit;"
elif [[ $action = "unshare" ]]
then
    paramString="selfservice bookmark select $bookmarkName unshare; commit;"
else
    usage
    exit 1
fi
result=$((${SSH_CMD} $paramString)
check_result
if [[ $action = "share" ]]
then
    echo "Bookmark ${bookmarkName} is now in shared mode"
elif [[ $action = "unshare" ]]
then
    echo "Bookmark ${bookmarkName} is now in not-share mode"
fi
}

##### Main

while getopts "u:d:a:h" flag; do
    case "$flag" in
        u )          username=${OPTARG%:*}
                    ;;
        d )          engine=$OPTARG
                    ;;
        a )          action=$OPTARG
                    ;;
        h )          usage
                    exit
                    ;;
        * )          usage
                    exit 1
    esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there is 1 positional arguments
if [ $# != 1 ]
then
    usage
    exit 1
fi

```

```
# Get the one positional arguments  
bookmarkName=$1  
  
create_session  
bookmark_action
```

CLI Cookbook: How to Update a Delphix Self-Service Bookmark

Prerequisites:

- Have Delphix Self-Service user privileges
- Know the bookmark you would like to update

Procedure:

1. Log into Delphix Engine as a Delphix Self-Service user or admin.

```
ssh jsUser@<yourengine>
```

2. Navigate to the Delphix Self-Service bookmarks, and choose the one you would like to update.

```
jsUser > selfservice bookmark  
jsUser selfservice bookmark > ls  
jsUser selfservice bookmark > select <jsBookmark>
```

3. Update the bookmark.

```
jsUser selfservice bookmark <jsBookmark> *> update  
jsUser selfservice bookmark <jsBookmark> update *> set tags="tag text"  
jsUser selfservice bookmark <jsBookmark> update *> commit
```

CLI Cookbook: Hooks and Hook Templates

These topics describe the command-line interface procedures for working with hooks and hook templates.

This section covers the following topics:

- [CLI Cookbook: Changing the PowerShell version associated with a hook template](#)

CLI Cookbook: Changing the PowerShell version associated with a hook template

This article is to document how to change the PowerShell version associated with a hook template using the CLI.

 This operation is not allowed through GUI.

Procedure:

1. Login into Delphix Engine CLI using "admin" credentials.

```
ssh admin@<delphix_engine>
```

2. Go to **source** and then to **operationTemplate**.

```
delphix> source
delphix source> operationTemplate
```

3. Execute the **ls** command to see the existing hooks template and select one of the existing hook templates which you want to update.

```
delphix source operationTemplate> ls
Objects
NAME                DESCRIPTION                LASTUPDATED
OPERATION.TYPE
Test_Template       Test_Template              2020-06-12T04:15:01.759Z
RunDefaultPowerShellOnSourceOperation

delphix source operationTemplate> select Test_Template

delphix source operationTemplate 'Test_Template'> ls
Properties
  type: OperationTemplate
  name: Test_Template
  description: Test_Template
  lastUpdated: 2020-06-12T04:15:01.759Z
  operation:
    type: RunDefaultPowerShellOnSourceOperation
    name: Test_Template
    command: echo_command
    reference: OPERATION_TEMPLATE-1
```

4. Run the **update** command and change the hook template type from **RunDefaultPowerShellOnSourceOperation** to **RunPowerShellOnSourceOperation**.

 Users can also change the template type from *RunPowerShellOnSourceOperation* to *RunDefaultPowerShellOnSourceOperation*.

```
delphix source operationTemplate 'Test_Template'> update
```

```
delphix source operationTemplate 'Test_Template' update *> set
operation.type=RunPowerShellOnSourceOperation
delphix source operationTemplate 'Test_Template' update *> ls
Properties
  type: OperationTemplate
  name: Test_Template
  description: Test_Template
  operation:
    type: RunPowerShellOnSourceOperation (*)
    name: Test_Template
    command: echo_command
```

5. Run the **commit** command to perform the update on the hook template. Execute the **ls** command again to see the change in the hook template type.

```
delphix source operationTemplate 'Test_Template' update *> commit

delphix source operationTemplate 'Test_Template'> ls
Properties
  type: OperationTemplate
  name: Test_Template
  description: Test_Template
  lastUpdated: 2020-06-12T04:30:44.203Z
  operation:
    type: RunPowerShellOnSourceOperation
    name: Test_Template
    command: echo_command
  reference: OPERATION_TEMPLATE-1
```



1. Template type can only be changed from "RunDefaultPowerShellOnSourceOperation" to "RunPowerShellOnSourceOperation" or vice versa. You **can not** change the type to other available types. If you try to do that, you will get this error: "Error: The type of this operation (RUN_POWERSHELL_ON_SOURCE_OPERATION) cannot be modified after creation.".
2. Template type "RunPowerShellOnSourceOperation" denotes "PowerShell version 2" and "RunDefaultPowerShellOnSourceOperation" denotes "Default PowerShell version". Here, default version is the version of PowerShell installed on the target host.

CLI Cookbook: Network Performance

These topics describe command-line interface procedures using the iperf-based Network Performance Tool:

- [CLI Cookbook: Delphix Session Protocol Test from Primary Engine to Replication Engine](#)
- [CLI Cookbook: Running the Network Test via CLI - Latency](#)
- [CLI Cookbook: Running the Network Test via CLI - Throughput](#)

CLI Cookbook: Delphix Session Protocol Test from Primary Engine to Replication Engine

Prerequisites

The network performance tool measures network performance between a Delphix Engine and an environment host. You must have added an environment in order to use this tool.

This transmission control protocol (TCP) throughput test uses TCP port 50001 by default. The port can also be configured on a per-test-run basis. For the duration of a given throughput test, a server on the receiver will be listening on this port. For a transmit test, the receiver is the remote host; for a receive test, the receiver is the Delphix Engine

Procedure

Delphix uses the Delphix Session Protocol (DSP) protocol to communicate between primary and replication engines.

1. Login to Delphix Engine using a Delphix administrator account such as **delphix_admin**. As soon as you login, you will get prompt with engine name.

```
-bash-4.3$ ssh delphix_admin@delphix
Password:
delphix>
```

2. Create a DSP test.

```
delphix> network test dsp create
```

3. Set the destinationType to DELPHIX_ENGINE to do a DSP test between two Delphix Engines. The default is REMOTE_HOST which executes a DSP test between a source or target host and the Delphix Engine.

```
delphix network test dsp create *> set destinationType=DELPHIX_ENGINE
```

4. Use **get** to see other optional arguments. Modify the test parameters as needed and commit to start the test.

```
delphix network test dsp create *> get
type: NetworkDSPTestParameters
blockSize: 64KB
compression: false
destinationType: REMOTE_HOST
direction: TRANSMIT
duration: 30
encryption: false
numConnections: 0
queueDepth: 32
receiveSocketBuffer: 256KB
remoteDelphixEngineInfo: (unset)
remoteHost: (unset)
sendSocketBuffer: 256KB
```

```

delphix network test dsp create *> set remoteDelphixEngineInfo.address=delphix2
delphix network test dsp create *> set
remoteDelphixEngineInfo.principal=delphix_admin
delphix network test dsp create *> edit remoteDelphixEngineInfo
delphix network test dsp create remoteDelphixEngineInfo *> get
  type: RemoteDelphixEngineInfo (*)
  address: delphix2 (*)
  credential: (required)
  principal: delphix_admin (*)
delphix network test dsp create remoteDelphixEngineInfo *> set
credential.password=delphix
delphix network test dsp create remoteDelphixEngineInfo *> commit
`NETWORK_DSP_TEST-2
Dispatched job JOB-8
NETWORK_DSP_TEST_EXECUTE job started.
Measuring throughput with variable number of connections: 716 Mbps measured f
or 1 connections.
Measuring throughput with variable number of connections: 711 Mbps measured f
or 2 connections.
Measuring throughput with variable number of connections: 611 Mbps measured f
or 4 connections.
Measuring throughput with variable number of connections: 646 Mbps measured f
or 6 connections.
Measuring throughput with variable number of connections: 567 Mbps measured f
or 8 connections.
Measuring average throughput for 30 seconds with 1 connections.
Measured throughput of 408 Mbps.
NETWORK_DSP_TEST_EXECUTE job completed successfully.

```

5. Retrieve the test results.

```

delphix> network test dsp list
NAME                                PARAMETERS.DIRECTION  STATE      THROUGHPUT
delphix2-2018-01-23T21:25:31.172Z  TRANSMIT              COMPLETED 880.5Mbps
delphix2-2018-02-02T17:54:58.322Z  TRANSMIT              COMPLETED 408.5Mbps

```

Related Topics

- [Network Performance Test Tool Interface](#)
- [CLI Cookbook: Running the Network Test via CLI - Throughput](#)
- [CLI Cookbook: Running the Network Test via CLI - Latency](#)

CLI Cookbook: Running the Network Test via CLI - Latency

Prerequisites

The network performance tool measures network performance between a Delphix Engine and an environment host. You must have added an environment in order to use this tool.

This transmission control protocol (TCP) throughput test uses TCP port 50001 by default. The port can also be configured on a per-test-run basis. For the duration of a given throughput test, a server on the receiver will be listening on this port. For a transmit test, the receiver is the remote host; for a receive test, the receiver is the Delphix Engine.

Procedure

The network latency test measures network round-trip latency by transmitting ICMP echo requests (like the ping utility) and measuring the time to receive replies from the remote host. To execute a test:

1. Login as a domain user to the Delphix Engine CLI using ssh.
2. Create a network latency test.

```
delphix> network test latency
delphix network test latency> create
delphix network test latency create *>
```

3. You must set **remoteHost** to the name of an environment host already configured in the Delphix Engine. (You should press tab after the "=" (equal sign) to auto-populate and confirm registered destinations). Use 'get' to see other optional arguments. Modify the test parameters as needed and **commit** to start the test.

```
delphix network test latency create *> set remoteHost=oracletarget
delphix network test latency create *> get
  type: NetworkLatencyTestExecuteParameters
  remoteHost: oracletarget
  requestCount: 20
  requestSize: 8B
delphix network test latency create *> commit
  Dispatched job JOB-20
  NETWORK_LATENCY_TEST_EXECUTE job started for
  "oracletarget-2014-06-20T18:57:28.659Z".
  Executing network latency test.
  NETWORK_LATENCY_TEST_EXECUTE job for
  "oracletarget-2014-06-20T18:57:28.659Z" completed successfully.
```

4. The job will be submitted and visible in the Delphix Management application.
5. Retrieve the test results. All times are in microseconds.

```
delphix network test latency> list
NAME                                AVERAGE
oraclesource-2014-06-20T18:57:28.659Z  872
delphix network test latency> select oraclesource-2014-06-20T18:57:28.659Z
delphix network test latency "oraclesource-2014-06-20T18:57:28.659Z"> get
  type: NetworkLatencyTest
```

```
name: oraclesource-2014-06-20T18:57:28.659Z
average: 872
endTime: 2014-06-20T18:57:48.558Z
loss: 0
maximum: 2755
minimum: 294
reference: NETWORK_LATENCY_TEST-2
remoteAddress: 172.16.203.184
remoteHost: oraclesource
requestCount: 20
requestSize: 8B
startTime: 2014-06-20T18:57:28.659Z
stddev: 527
```

Related Topics

- [Network Performance Test Tool Interface](#)
- [CLI Cookbook: Running the Network Test via CLI - Throughput](#)
- [CLI Cookbook: Delphix Session Protocol Test from Primary Engine to Replication Engine](#)

CLI Cookbook: Running the Network Test via CLI - Throughput

Prerequisites

The network performance tool measures network performance between a Delphix Engine and an environment host. You must have added an environment in order to use this tool.

This transmission control protocol (TCP) throughput test uses TCP port 50001 by default. The port can also be configured on a per-test-run basis. For the duration of a given throughput test, a server on the receiver will be listening on this port. For a transmit test, the receiver is the remote host; for a receive test, the receiver is the Delphix Engine.

Procedure

The network throughput test measures sustained throughput using a synthetic workload to or from a remote host. To execute a test:

1. Login as a domain user to the **Delphix Engine CLI** using ssh.
2. Create a network throughput test.

```
delphix> network test throughput
delphix network test throughput> create
delphix network test throughput create *>
```

3. You must set **remoteHost** to the name of an environment host already configured in the Delphix Engine. Use 'get' to see other optional arguments. Modify the test parameters as needed and **commit** to start the test.

```
delphix network test throughput create *> set remoteHost=oraclesource
delphix network test throughput create *> ls
Properties
  type: NetworkThroughputTestParameters
  blockSize: 128KB
  direction: TRANSMIT
  duration: 30
  numConnections: 0
  port: 50001
  receiveSocketBuffer: 4MB
  remoteHost: oraclesource
  sendSocketBuffer: 4MB
delphix network test throughput create *> commit
  Dispatched job JOB-21
  NETWORK_THROUGHPUT_TEST_EXECUTE job started for
  "oraclesource-2014-06-20T19:30:12.566Z".
  Executing network throughput transmit test.
  Measuring throughput with variable number of connections: 1.
  Measuring throughput with variable number of connections: 2.
  Measuring throughput with variable number of connections: 4.
  Measuring throughput with variable number of connections: 6.
  Measuring throughput with variable number of connections: 8.
  Measuring maximum sustained throughput for 30 seconds with 8 connections.
  NETWORK_THROUGHPUT_TEST_EXECUTE job for
  "oraclesource-2014-06-20T19:30:12.566Z" completed successfully.
```

4. The job will be submitted and visible in the Delphix Management application.
Retrieve the test results.

```
delphix network test throughput> list
NAME                                DIRECTION  STATE      THROUGHPUT
oraclesource-2014-06-20T19:30:12.566Z  TRANSMIT   COMPLETED 695.6Mbps
delphix network test throughput> select oraclesource-2014-06-20T19:30:12.566Z
delphix network test throughput "oraclesource-2014-06-20T19:30:12.566Z"> get
  type: NetworkThroughputTest
  name: oraclesource-2014-06-20T19:30:12.566Z
  endTime: 2014-06-20T19:31:15.041Z
  numConnections: 8
  parameters:
    type: NetworkThroughputTestParameters
    blockSize: 128KB
    direction: TRANSMIT
    duration: 30
    numConnections: 0
    port: 50001
    receiveSocketBuffer: 4MB
    remoteHost: oraclesource
    sendSocketBuffer: 4MB
  reference: NETWORK_THROUGHPUT_TEST-2
  remoteAddress: 172.16.203.184
  startTime: 2014-06-20T19:30:12.566Z
  state: COMPLETED
  throughput: 695.6Mbps
```

Related Topics

- [Network Performance Test Tool Interface](#)
- [CLI Cookbook: Running the Network Test via CLI - Latency](#)
- [CLI Cookbook: Delphix Session Protocol Test from Primary Engine to Replication Engine](#)

Kerberos CLIs

The following topics provide information related to the Kerberos CLIs.

- [CLI Cookbook: Admin Application Configuration](#)
- [CLI Cookbook System app configuration](#)

CLI Cookbook: Admin Application Configuration

The following are done after logging into the Delphix admin app as `admin@<delphix engine hostname>`.

- [Add Kerberos Environment](#)
- [Add Kerberos ASE Instance](#)
- [Link a dSource](#)

Add Kerberos Environment

```
jkb-5160.dcenter> cd /environment/  
jkb-5160.dcenter environment> ls  
  
Children  
  
oracle  
  
user  
  
windows  
  
Operations  
  
create  
  
jkb-5160.dcenter environment> create  
jkb-5160.dcenter environment create *> ls  
  
Properties  
  
  type: HostEnvironmentCreateParameters  
  
  hostEnvironment:  
    type: UnixHostEnvironment  
    name: (unset)  
    aseHostEnvironmentParameters: (unset)  
    description: (unset)  
  
  hostParameters:  
    type: UnixHostCreateParameters
```

```

host:
  type: UnixHost
  address: (required)
  privilegeElevationProfile: (unset)
  sshPort: 22
  toolkitPath: (required)
  truststorePassword: (unset)

```

```
primaryUser:
```

```

  type: EnvironmentUser
  name: (unset)
  credential:
    type: PasswordCredential
    password: (required)
  environment: (unset)
  groupId: (unset)
  userId: (unset)

```

```
jkb-5160.dcenter environment create *> set hostParameters.host.address=ln-rh64-tgt.dc2.delphix.com
```

```
jkb-5160.dcenter environment create *> set hostParameters.host.toolkitPath=/tmp
```

```
jkb-5160.dcenter environment create *> set
primaryUser.credential.type=KerberosCredential
```

```
jkb-5160.dcenter environment create *> commit
```

```
`UNIX_HOST_ENVIRONMENT-2
```

```
Dispatched job JOB-4
```

```
ENVIRONMENT_CREATE_AND_DISCOVER job started for "ln-rh64-tgt.dc2.delphix.com".
```

```
ENVIRONMENT_CREATE_AND_DISCOVER job for "ln-rh64-tgt.dc2.delphix.com" completed
successfully.
```

```
jkb-5160.dcenter environment>
```

```
jkb-5160-2.dcenter environment> ls
```

Objects

NAME	DESCRIPTION
------	-------------

ln-rh64-tgt.dc2.delphix.com	-
-----------------------------	---

Children

oracle

user

windows

Operations

Create

```
jkb-5160-2.dcenter environment> select ln-rh64-tgt.dc2.delphix.com
```

```
jkb-5160-2.dcenter environment 'ln-rh64-tgt.dc2.delphix.com'> ls
```

Properties

type: UnixHostEnvironment

name: ln-rh64-tgt.dc2.delphix.com

aseHostEnvironmentParameters: (unset)

description: (unset)

enabled: **true**

host: ln-rh64-tgt.dc2.delphix.com

primaryUser: sybase

reference: UNIX_HOST_ENVIRONMENT-1

Operations

delete

update

```
disable  
enable  
refresh  
jkb-5160-2.dcenter environment 'ln-rh64-tgt.dc2.delphix.com'>
```

Add Kerberos ASE Instance

```
jkb-5160-2.dcenter environment> cd /repository/  
jkb-5160-2.dcenter repository> ls  
  
Children  
  
template  
  
Operations  
  
create  
  
compatibleRepositories  
  
jkb-5160-2.dcenter repository> create  
jkb-5160-2.dcenter repository create *> ls  
  
Properties  
  
  type: ASEInstance  
  
  credentials: (unset)  
  
  dbUser: (unset)  
  
  environment: (required)  
  
  installationPath: (required)  
  
  instanceName: (required)  
  
  instanceOwner: (required)  
  
  ports: (required)  
  
  servicePrincipalName: (unset)  
  
  version: (unset)
```

```

jkb-5160-2.dcenter repository create *> set credentials.type=KerberosCredential
jkb-5160-2.dcenter repository create *> set environment=ln-rh64-tgt.dc2.delphix.com
jkb-5160-2.dcenter repository create *> set installationPath=/opt/sybase/15-0
jkb-5160-2.dcenter repository create *> set ports=5100
jkb-5160-2.dcenter repository create *> set instanceName=ASE1570_S1
jkb-5160-2.dcenter repository create *> set instanceOwner=sybase
jkb-5160-2.dcenter repository create *> set servicePrincipalName=ASE1570_S1
jkb-5160-2.dcenter repository create *> commit

  `ASE_INSTANCE-1
jkb-5160-2.dcenter repository> ls

Objects

NAME          VERSION ENVIRONMENT
ASE1570_S1    15.7 SP138 ln-rh64-tgt.dc2.delphix.com

Children

template

Operations

create

compatibleRepositories

jkb-5160-2.dcenter repository> select ASE1570_S1
jkb-5160-2.dcenter repository 'ASE1570_S1'> ls

Properties

  type: ASEInstance

  name: ASE1570_S1

  credentials:

    type: KerberosCredential

```

```
dbUser: sybase
discovered: false
environment: ln-rh64-tgt.dc2.delphix.com
installationPath: /opt/sybase/15-0
instanceName: ASE1570_S1
instanceOwner: sybase
instanceOwnerGid: 500
instanceOwnerUid: 500
internalVersion: 15.7 SP138
linkingEnabled: true
pageSize: 4096
ports: 5100
provisioningEnabled: true
reference: ASE_INSTANCE-1
servicePrincipalName: ASE1570_S1
staging: false
version: 15.7 SP138
```

Operations

delete

update

```
jkb-5160-2.dcenter repository 'ASE1570_S1'>
```

Link a dSource

```
jkb-5160-2.dcenter> cd /database
```

```
jkb-5160-2.dcenter database> ls
```

Children

template

Operations

createEmpty

createRestorationDataset

export

fileMapping

link

oracleSupportedCharacterSets

provision

validateXpp

xpp

```
jkb-5160-2.dcenter database> link
```

```
jkb-5160-2.dcenter database link *> ls
```

Properties

type: LinkParameters

name: (required)

description: (unset)

group: (required)

linkData:

type: ASELinkData

config: (required)

dbCredentials:

type: PasswordCredential

password: (required)

dbUser: (unset)

dumpCredentials: (unset)

```

externalFilePath: (unset)
loadBackupPath: (required)
loadLocation: (unset)
operations: (unset)
sourceHostUser: (required)
sourcingPolicy: (unset)
stagingHostUser: (required)
stagingPostScript: (unset)
stagingPreScript: (unset)
stagingRepository: (required)
syncParameters:
    type: ASELatestBackupSyncParameters
validatedSyncMode: ENABLED

```

Operations

defaults

```

jkb-5160-2.dcenter database link *> set name=test1
jkb-5160-2.dcenter database link *> set group=Untitled
jkb-5160-2.dcenter database link *> set linkData.config=test1
jkb-5160-2.dcenter database link *> set
linkData.dbCredentials.type=KerberosCredential
jkb-5160-2.dcenter database link *> set linkData.loadBackupPath=/home/sybase/db
jkb-5160-2.dcenter database link *> set linkData.stagingRepository=ASE1570_S1
jkb-5160-2.dcenter database link *> set linkData.sourceHostUser=sybase
jkb-5160-2.dcenter database link *> set linkData.stagingHostUser=sybase
jkb-5160-2.dcenter database link *> set
linkData.syncParameters.type=ASENewBackupSyncParameters

```

```
jkb-5160-3.dcenter database link linkData syncParameters *> commit
```

```
  `ASE_DB_CONTAINER-1
```

```
  Dispatched job JOB-4
```

```
  DB_LINK job started for "Untitled/test1".
```

```
  DB_LINK job for "Untitled/test1" completed successfully.
```

```
jkb-5160-2.dcenter database>
```

```
jkb-5160-2.dcenter database> ls
```

```
Objects
```

```
NAME    PROVISIONCONTAINER  DESCRIPTION
```

```
test1   -                   -
```

```
Children
```

```
template
```

```
Operations
```

```
createEmpty
```

```
createRestorationDataset
```

```
export
```

```
fileMapping
```

```
link
```

```
oracleSupportedCharacterSets
```

```
provision
```

```
validateXpp
```

```
xpp
```

```
jkb-5160-2.dcenter database>
```

CLI Cookbook System app configuration

The following are done after logging into the Delphix System configuration application as `sysadmin@<delphix_engine_hostname>`.

Define the system Kerberos configuration

```
jkb-5160.dcenter> cd /service/kerberos/
jkb-5160.dcenter service kerberos> ls

Properties

  type: KerberosConfig

  enabled: false

  kdcs:

    0:

      type: KerberosKDC

      hostname:

      port: 88

  keytab: (unset)

  principal:

  realm:

Operations

update

reset

jkb-5160.dcenter service kerberos> update

jkb-5160.dcenter service kerberos update *> set realm=DELPHIX.COM

jkb-5160.dcenter service kerberos update *> set principal=sybase

jkb-5160.dcenter service kerberos update *> set
keytab=BQIAAABEAAEAC0RFTFBISVguQ09NAAZzeWJhc2UAAAABWNVqQgMAEgAgjF/
zgNuw27Uy9vEgvPvpeevAaAw6c5HaVAWVSLZnmngAAAA0AAEAC0RFTFBISVguQ09NAAZzeWJhc2UAAAABWNVq
QgMAEQAQTsNZOXmHhaeuXPeyQRXeLAAAADwAAQALREVMUEhJWC5DT00ABnN5YmFzZQAAAAFY1WpCAwAQABjHS
```

```
VHfC+p8yIYfgFTIv1T7m99n2Spn7wIAAAA0AAEAC0RFTFBISVguQ09NAAZzeWJhc2UAAAABWNVqQgMAFwAQhJ
PmRokSfz310fe3eVUD3wAAACwAAALREVMUEhJWC5DT00ABnN5YmFzZQAAAAFY1WpCAwAIAAhiFUmfjttmWA
AACwAAALREVMUEhJWC5DT00ABnN5YmFzZQAAAAFY1WpCAwADAAj4rgff+Au/ng==
```

```
jkb-5160.dcenter service kerberos update *> edit kdc
```

```
jkb-5160.dcenter service kerberos update kdc *> edit 0
```

```
jkb-5160.dcenter service kerberos update kdc 0 *> set hostname=kerberos-01.delphix.com
```

```
jkb-5160.dcenter service kerberos update kdc 0 *> back
```

```
jkb-5160.dcenter service kerberos update kdc *> back
```

```
jkb-5160.dcenter service kerberos update *> ls
```

Properties

```
type: KerberosConfig
```

```
kdc:
```

```
0:
```

```
type: KerberosKDC (*)
```

```
hostname: kerberos-01.delphix.com (*)
```

```
port: 88 (*)
```

```
keytab: ***** (*)
```

```
principal: sybase (*)
```

```
realm: DELPHIX.COM (*)
```

```
jkb-5160.dcenter service kerberos update *> commit
```

```
jkb-5160.dcenter service kerberos> ls
```

Properties

```
type: KerberosConfig
```

```
enabled: true
```

```
kdc:
```

```
0:
```

```
type: KerberosKDC
```

```
hostname: kerberos-01.delphix.com
port: 88
keytab: (unset)
principal: sybase
realm: DELPHIX.COM
```

Operations

update

reset

```
jkb-5160.dcenter service kerberos>
```

Clear Kerberos Configuration

```
jkb-5160.dcenter service kerberos> reset
jkb-5160.dcenter service kerberos reset *> commit
jkb-5160.dcenter service kerberos> ls
```

Properties

```
type: KerberosConfig
enabled: false
kdc:
  0:
    type: KerberosKDC
    hostname:
    port: 88
keytab: (unset)
principal:
realm:
```

Operations

update

reset

jkb-5160.dcenter service kerberos>

Web Services API Guide

These topics describe interfacing with the public web service APIs, building automation facilities and integrating with third-party orchestration tools.

This section covers the following topics:

- [API Version Information](#)
- [Web Service Object Model](#)
- [Web Service Protocol](#)
- [CLI to Web Services Translation](#)
- [GUI API Mapping](#)
- [CLI to Python Transition](#)
- [Python Cookbook: Adding a UNIX Host](#)
- [So You Want to Work with Delphix APIs?](#)
- [API Cookbook: Common Tasks, Workflows, and Examples](#)
- [Kerberos APIs](#)

API Version Information

This topic describes API version information for each release of the Delphix Engine, including schema changes and links to the relevant version of the schema.

Delphix Engine Version	API Version	Link to Schema within the Appliance
5.0.0.x - 5.0.3.x	1.7.0	<a href="http://<engine-address>/api/json/versions/1.7.0/delphix.json">http://<engine-address>/api/json/versions/1.7.0/delphix.json
5.0.4.x	1.7.1	<a href="http://<engine-address>/api/json/versions/1.7.1/delphix.json">http://<engine-address>/api/json/versions/1.7.1/delphix.json
5.1.0.x - 5.1.2.x	1.8.0	<a href="http://<engine-address>/api/json/versions/1.8.0/delphix.json">http://<engine-address>/api/json/versions/1.8.0/delphix.json
5.1.3.x - 5.1.5.x	1.8.1	<a href="http://<engine-address>/api/json/versions/1.8.1/delphix.json">http://<engine-address>/api/json/versions/1.8.1/delphix.json
5.1.6.0	1.8.2	<a href="http://<engine-address>/api/json/versions/1.8.2/delphix.json">http://<engine-address>/api/json/versions/1.8.2/delphix.json
5.2.0.0 - 5.2.2.1	1.9.0	<a href="http://<engine-address>/api/json/versions/1.9.0/delphix.json">http://<engine-address>/api/json/versions/1.9.0/delphix.json
5.2.3.0 - 5.2.3.1	1.9.1	<a href="http://<engine-address>/api/json/versions/1.9.1/delphix.json">http://<engine-address>/api/json/versions/1.9.1/delphix.json
5.2.4.0	1.9.2	<a href="http://<engine-address>/api/json/versions/1.9.2/delphix.json">http://<engine-address>/api/json/versions/1.9.2/delphix.json
5.2.5.0 - 5.2.6.2	1.9.3	<a href="http://<engine-address>/api/json/versions/1.9.3/delphix.json">http://<engine-address>/api/json/versions/1.9.3/delphix.json
5.3.0.0 - 5.3.0.3	1.10.0	<a href="http://<engine-address>/api/json/versions/1.10.0/delphix.json">http://<engine-address>/api/json/versions/1.10.0/delphix.json
5.3.1.0	1.10.1	<a href="http://<engine-address>/api/json/versions/1.10.1/delphix.json">http://<engine-address>/api/json/versions/1.10.1/delphix.json
5.3.2.0	1.10.2	<a href="http://<engine-address>/api/json/versions/1.10.2/delphix.json">http://<engine-address>/api/json/versions/1.10.2/delphix.json

Delphix Engine Version	API Version	Link to Schema within the Appliance
5.3.3.0 - 5.3.3.1	1.10.3	<a href="http://<engine-address>/api/json/versions/1.10.3/delphix.json">http://<engine-address>/api/json/versions/1.10.3/delphix.json
5.3.4.0	1.10.4	<a href="http://<engine-address>/api/json/versions/1.10.4/delphix.json">http://<engine-address>/api/json/versions/1.10.4/delphix.json
5.3.5.0	1.10.5	<a href="http://<engine-address>/api/json/versions/1.10.5/delphix.json">http://<engine-address>/api/json/versions/1.10.5/delphix.json
5.3.6.0-5.3.9.0	1.10.6	<a href="http://<engine-address>/api/json/versions/1.10.6/delphix.json">http://<engine-address>/api/json/versions/1.10.6/delphix.json
6.0.0.0	1.11.0	<a href="http://<engine-address>/api/json/versions/1.11.0/delphix.json">http://<engine-address>/api/json/versions/1.11.0/delphix.json
6.0.1.0	1.11.1	<a href="http://<engine-address>/api/json/versions/1.11.1/delphix.json">http://<engine-address>/api/json/versions/1.11.1/delphix.json
6.0.2.0	1.11.2	<a href="http://<engine-address>/api/json/versions/1.11.2/delphix.json">http://<engine-address>/api/json/versions/1.11.2/delphix.json
6.0.3.0	1.11.3	<a href="http://<engine-address>/api/json/versions/1.11.3/delphix.json">http://<engine-address>/api/json/versions/1.11.3/delphix.json
6.0.4.0	1.11.4	<a href="http://<engine-address>/api/json/versions/1.11.4/delphix.json">http://<engine-address>/api/json/versions/1.11.4/delphix.json
6.0.5.0	1.11.5	<a href="http://<engine-address>/api/json/versions/1.11.5/delphix.json">http://<engine-address>/api/json/versions/1.11.5/delphix.json
6.0.6.0	1.11.6	<a href="http://<engine-address>/api/json/versions/1.11.6/delphix.json">http://<engine-address>/api/json/versions/1.11.6/delphix.json
6.0.7.0	1.11.7	<a href="http://<engine-address>/api/json/versions/1.11.7/delphix.json">http://<engine-address>/api/json/versions/1.11.7/delphix.json
6.0.8.0	1.11.8	<a href="http://<engine-address>/api/json/versions/1.11.8/delphix.json">http://<engine-address>/api/json/versions/1.11.8/delphix.json
6.0.9.0	1.11.9	<a href="http://<engine-address>/api/json/versions/1.11.9/delphix.json">http://<engine-address>/api/json/versions/1.11.9/delphix.json

Delphix Engine Version	API Version	Link to Schema within the Appliance
6.0.10.0	1.11.10	<a href="http://<engine-address>/api/json/versions/1.11.10/delphix.json">http://<engine-address>/api/json/versions/1.11.10/delphix.json
6.0.11.0	1.11.11	<a href="http://<engine-address>/api/json/versions/1.11.11/delphix.json">http://<engine-address>/api/json/versions/1.11.11/delphix.json
6.0.12.0	1.11.12	<a href="http://<engine-address>/api/json/versions/1.11.12/delphix.json">http://<engine-address>/api/json/versions/1.11.12/delphix.json
6.0.13.0	1.11.13	<a href="http://<engine-address>/api/json/delphix.json">http://<engine-address>/api/json/delphix.json

For more details see [API Changes](#).

Web Service Object Model

This topic describes the Delphix object model as exported over the web services.

Object Types

All objects in the Delphix API are "typed objects." All such objects have a `type` field that indicates the type of the object and its associated semantics. This allows for object inheritance and polymorphism without requiring separate APIs for each type and allows generic client-specific semantic encoding and decoding without having to be aware of the context. This means that even APIs that operate only a specific type (such as the `Group` API) still require a type field to be specified as part of the input, and will continue to report the type of objects when listing or retrieving objects. This allows the APIs to evolve in a backward-compatible fashion through the introduction of new types.

Certain "root" object types (groups, containers, sources, etc) have an associated API. This API is rooted at a particular point under `/resources/json/delphix`, but all APIs follow a standard format beneath this namespace. The CLI namespace is a direct reflection of this URL, and the API reference has an index both by object type as well as by object (CLI) path. These APIs may operate on different extended types (such as `OracleSIConfig` and `OracleRACConfig`), but the base operations remain the same regardless of input type.

Object References

Some objects returned by the APIs are pure typed objects, in that they don't represent the persistent state on the Delphix Engine but are rather calculated and returned upon request. Most of the objects in the system, however, are "persistent objects." Persistent objects (of type `PersistentObject`) have a stable reference that uniquely identifies the object on the system. This reference is separate from its name so that objects can be renamed without affecting the programmatic API. More information about object names and how they can be represented generically can be found in the [CLI documentation](#). Object references are opaque tokens; while they can be stored and reused for later use, and interpretation of their contents is unstable and may break in a future release.

The Delphix object hierarchy is stitched together by these object references. When fetching an object that refers to another object, the member will be returned as a reference, rather than being inserted directly within the parent object. This allows consumers to control when and how links are resolved and make it clear when an object may change independently from its parent. The per-object APIs outlined below all operate on object references.

Note that some Delphix objects are `singleton` objects, and there is only one such object on the system. These objects do not have references because the API URL uniquely identifies the object on the system.

API Operations

All APIs optionally support the following operations:

- `CREATE` - Create a new instance of the given object type. This is used for most objects, but more complicated objects, such as `dSources` and `VDBs`, must be created through a dedicated `link` or `provision` operation. The input to this operation is typically the object itself, though some objects may have specialized parameters used during the creation of objects. An example of this is `HostEnvironmentCreateParameters`.
- `UPDATE` - Update properties of the given object, specified as an object reference in the URL.

- **DELETE** - Delete a particular object, specified as an object reference in the URL. These operations are typically done as HTTP **DELETE** operations, but it is also possible to do a **POST** operation to the `/delete` API to do the same thing. The **POST** form allows for delete-specific parameters, such as `OracleDeleteParameters`.
- **GET** - Get the properties for a particular object, specified as an object reference in the URL.
- **LIST** - List all objects of the given type within the system. These APIs typically take optional query parameters that allow the set of results to be constrained, filtered, paginated, or sorted.

In addition, the following non-CRUD operations may be supported:

- **Root Operation** - A **POST** or **GET** operation to the root of an API namespace, not associated with a particular object. This can be used for singleton objects, such as `NDMPConfig`, operations that create objects, such as `Link`, and operations that operate on multiple objects at once.
- **Per-object Operation** - A **POST** operation to a particular object reference. These operations are typically read-write but are not required to be so. These would include read-only operations that cannot be modeled as CRUD operations or require complex input use per-object operations.

Database Object Model

In order to support a wide variety of databases and database configurations, the database object model is more complex than it may initially appear after having used the Delphix Management application. For example, there is no such thing as a "dSource" or "VDB" object, only data "containers" with attached "sources". More information about how Database objects are modeled within Delphix can be found in the [CLI documentation](#).

Asynchronous Execution

All APIs are designed to be transactionally safe and "quick." However, there are operations that may take a long period of time or may need to reach out to external hosts or databases such that they cannot be done safely within the context of a single API call. Such operations will dispatch a `Job` to handle asynchronous execution of the operation. **Any API** can potentially spawn a job, and which APIs spawn jobs and which do not may differ between object types or releases. If you are developing a full-featured automation system, it is recommended that you build a generic infrastructure to handle job monitoring, rather than encoding the behavior of particular APIs that may change over time.

Every operation, except for `LIST` and `GET`, which are guaranteed to be read-only, can potentially spawn a job.

This is represented by the `job` field of the `APIResult` object. If this field is `null`, then the action can be completed within the bounds of the API call. Otherwise, a reference to a dispatched job is returned.

Jobs can spawn other jobs for especially complex operations, such as provisioning to an Oracle cluster environment. The job returned in the API invocation is the root job, and overall success or failure of the operation is determined by the state of this job. Some operations may succeed even if a child job fails. An example would be provisioning to an Oracle cluster where one node failed, but others were successful.

Progress can be monitored by examining the `JobEvent` objects in the `Job` object returned through the job API.

Web Service Protocol

This topic describes an overview of the web service API and available facilities.

Introduction

The Delphix Engine provides a set of public stable web service APIs (application programming interfaces). The web services form the basis upon which the GUI and CLI are built, and are designed to be **public** and **stable**. This guide covers the basic operation of the protocol, concepts, and considerations when building layered infrastructure. It is not a reference for all available APIs. For more information on available APIs, go to the '/api' URL of a Delphix appliance, which will provide a complete reference of all available APIs for the version of Delphix running on that system.

```
http://<server>/api
```

The CLI is a thin veneer over the web services. If you are new to the web services, it is recommended you first test out operations with the CLI, and use the `setopt trace=true` option to dump the raw data being sent and received to see the API in action.

Protocol Operation

The Delphix web services are a [RESTful](#) API with loose [CRUD](#) semantics using [JSON](#) encoding.

The following HTTP methods are supported:

- `GET` - Retrieve data from the server where complex input is not needed. All `GET` requests are guaranteed to be read-only, but not all read-only requests are required to use `GET`. Simple input (strings, number, boolean values) can be passed as query parameters.
- `POST` - Issue a read/write operation, or make a read-only call that requires complex input. The optional body of the call is expressed as JSON.
- `DELETE` - Delete an object on the system. For languages that don't provide a native wrapper for `DELETE`, or for delete operations with optional input, all delete operations can also be invoked as `POST` to the same URL with `/delete` appended to it.

Regardless of the operation, the result is returned as a JSON encoded result, the contents of which are covered below. For example, the following invocation of `curl` demonstrates establishing a new Session (pretty-printing the result):

```
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/session \
-c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 4,
    "micro": 3
  }
}
```

```

EOF
{
  "status": "OK",
  "result": {
    "type": "APISession",
    "version": {
      "type": "APIVersion",
      "major": 1,
      "minor": 4,
      "micro": 3
    },
    "locale": "en_US",
    "client": null
  },
  "job": null
}
EOF

```

NOTE: It is generally recommended to set the API session version to the [highest level supported](#) by your Delphix Engine.

Session Establishment

Login involves establishing a session and then authenticating to the Delphix Engine. Only authenticated users can access web APIs. Each user must establish a session prior to making any other API calls. This is done by sending a `Session` object to the URL `/resources/json/delphix/session`. This session object will specify the `APIVersion` to use for communication between the client and server. If the server doesn't support the version requested due to an incompatible change reflected in the API major version number, an error will be returned.

The resulting session object encodes the native server version, which can be different than the version requested by the client. If the server is running a more recent but compatible version, any translation of input and output to the native version is handled automatically. More information on versioning can be found in the documentation for the `APIVersion` object within the API reference on a Delphix system. If the client supports multiple versions, the appropriate type can be negotiated by trying to establish a session with each major version supported and then inspecting the version returned.

The session will also return an identifier through browser cookies that can be reused in subsequent calls to use the same session credentials and state without having to re-authenticate. The format of this cookie is private to the server and may change at any point. Sessions do not persist across a server restart or reboot. The mechanism by which this cookie is preserved and sent with subsequent requests is client-specific. The following demonstrates invoking the session login API call using `curl` and storing the cookies in the file `~/cookies.txt` for later use:

```

$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 4,
    "micro": 3
  }
}

```

```

    }
  }
EOF
{
  "status": "OK",
  "result": {
    "type": "APISession",
    "version": {
      "type": "APIVersion",
      "major": 1,
      "minor": 4,
      "micro": 3
    },
    "locale": "en_US",
    "client": null
  },
  "job": null
}
EOF

```

Authentication

Once the session has been established, the next step is to authenticate to the server by executing the `LoginRequest` API. Unauthenticated sessions are prohibited from making any API calls other than this login request. The username can be either a system user or domain user, and the backend will authenticate using the appropriate method. This example illustrates logging in via curl using cookies created when the session was established:

```

$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/login \
-b cookies.txt -c cookies2.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "delphix_user",
  "password": "delphix_pass",
  "target": "DOMAIN"
}
EOF

```

The new cookie (cookie2.txt) will need to be used in subsequent API requests. The login API currently only supports authentication by a password. There is no way to authenticate using any shared key or alternate authentication strategy.

CLI to Web Services Translation

This topic describes using the CLI to understand public web service APIs.

The [command-line interface](#) is a direct translation of the web services API to an interactive environment. This allows you to use the CLI to explore functionality with tab completion, integrated help, stronger type checking, and an indication of expected types and required fields. When trying to determine how to invoke an operation through the web services or interpret the results, it is recommended that you first learn how to do the same through the CLI, and then use the provided tools to translate that into web services call.

CLI Translation to HTTP

The CLI namespace is identical to the web service URLs for each base object and operation type. The root of all web services is `/resources/json/delphix`. Any additional CLI context is appended to this URL, joined by slashes. For example:

```
delphix> database provision
```

Is equivalent to:

```
POST /resources/json/delphix/database/provision
```

All operations in the CLI (those that require an explicit `commit` command) are modeled as `POST` HTTP calls. This is an example of a "root operation", as they are invoked not on any particular object, but across the system as a whole. Operations that are invoked on a particular object use a URL specific to that object:

```
delphix> database "dexample" refresh
```

Is equivalent to:

```
POST /resources/json/delphix/database/ORACLE_DB_CONTAINER-3/refresh
```

While the CLI uses names to refer to objects, at the API level we use references. Persistent objects (those with a permanent unique identity) have a `reference` field that is used in all cases to refer to the object. This allows references to remain constant even if objects are renamed.

For the standard CRUD (create, read, update, delete) operations, the mapping is slightly different:

CLI Operation	HTTP API
<code>database list</code>	<code>GET /resources/json/delphix/database</code>
<code>database create</code>	<code>POST /resources/json/delphix/database</code>

CLI Operation	HTTP API
<code>database "dexample" get</code>	<code>GET /resources/json/delphix/database/ <reference></code>
<code>database "dexample" update</code>	<code>POST /resources/json/delphix/database/ <reference></code>
<code>database "dexample" delete</code>	<code>DELETE /resources/json/delphix/database/ <reference></code> <code>POST /resources/json/delphix/database/ <reference>/delete</code>

The `DELETE` operation has an optional `POST` form that can take complex input for clients that don't support sending a payload as part of a `DELETE` operation.

Tracing HTTP Calls

The CLI also provides facilities to see the raw HTTP calls being made as part of any operation. To start with, viewing data in JSON format (`setopt format=json`) will provide an example of what the raw output looks like from the server. In its raw form, the CLI does not make any attempt to interpret the results, so references are displayed as references (and not names), and sizes are displayed as their raw numeric value.

This is helpful for scripting, but the CLI also has a mode to display the requests being sent to the server, the responses received, and the URLs used. To enable this mode, run `setopt trace=true` . Once you have determined how to do something through the CLI, you can use this mode as the basis for building direct integration with the raw HTTP APIs.

```
delphix group> setopt trace=true
delphix group> create
delphix group create *> set name=example
delphix group create *> set description="this is an example"
delphix group create *> commit
=== POST /resources/json/delphix/group ===
{
  "type": "Group",
  "name": "example",
  "description": "this is an example"
}
=== RESPONSE ===
{
  "status": "OK",
  "result": "GROUP-3",
  "action": "ACTION-37",
  "job": null
}
```

```

=== END ===
GROUP-3
delphix group> "example"
delphix group "example"> delete
=== GET /resources/json/delphix/group/GROUP-3 ===
=== RESPONSE ===
{
  "status": "OK",
  "result": {
    "type": "Group",
    "reference": "GROUP-3",
    "namespace": null,
    "name": "example",
    "description": "this is an example"
  },
  "action": null,
  "job": null
}
=== END ===
delphix group "example" delete *> commit
=== POST /resources/json/delphix/group/GROUP-3/delete ===
{}
=== RESPONSE ===
{
  "status": "OK",
  "result": "",
  "action": "ACTION-38",
  "job": null
}
=== END ===
delphix group>

```

When using trace mode within the context of a specific object, we refresh the contents of the object before executing each command. This results in the `GET` request before the `delete` command in the above example.

GUI API Mapping

This topic describes how to map from GUI operations to the corresponding CLI operation.

It is not always straightforward to convert from the visual layout of the GUI to the corresponding CLI operations. This topic outlines some common operations and indicates how they are represented in the CLI, web services, and the API documentation.

dSource Operations

GUI	CLI	API Topic	Input Object	Web Services
Link	<code>database link</code>	Container	LinkParameters	POST /resources/json/delphix/database/link
Show configuration	<code>database "name" get</code> <code>source "name" get</code>	Container Source	-	GET /resources/json/delphix/database/{ref} GET /resources/json/delphix/source/{ref}
Sync	<code>database "name" sync</code>	Container	SyncParameters	POST /resources/json/delphix/database/{ref}/sync
Update	<code>database "name" update</code>	Container	Container	POST /resources/json/delphix/database/{ref}
Delete	<code>database "name" delete</code>	Container	DeleteParameters	POST /resources/json/delphix/database/{ref}/delete

GUI	CLI	API Topic	Input Object	Web Services
Detach	database "name" detachSource	Container	DetachSourceParameters	POST /resources/ json/delphix/ database/{ref}/ detachSource
Attach	database "name" attachSource	Container	AttachSourceParameters	POST /resources/ json/delphix/ database/{ref}/ attachSource
Disable	source "name" disable	Source	SourceDisableParameters	POST /resources/ json/delphix/ source/{ref}/ disable
Enable	source "name" enable	Source	SourceEnableParameters	POST /resources/ json/delphix/ source/{ref}/ enable

VDB Operations

GUI	CLI	API Topic	Input Object	Web Services
Provision	database provision	Container	ProvisionParameters	POST /resources/ json/delphix/ database/provision
V2P	database export	Container	ExportParameters	POST /resources/ json/delphix/ database/export

GUI	CLI	API Topic	Input Object	Web Services
Refresh	database "name" refresh	Container	RefreshParameters	POST /resources/ json/delphix/ database/{ref}/ refresh
Snapshot	database "name" sync	Container	SyncParameters	POST /resources/ json/delphix/ database/{ref}/sync
Update	database "name" update	Container	Container	POST /resources/ json/delphix/ database/{ref}
Delete	database "name" delete	Container	DeleteParameters	POST /resources/ json/delphix/ database/{ref}/ delete
Start	source "name" start	Source	StartParameters	POST /resources/ json/delphix/ source/{ref}/start
Stop	source "name" stop	Source	StopParameters	POST /resources/ json/delphix/ source/{ref}/stop
Enable	source "name" enable	Source	SourceEnableParameters	POST /resources/ json/delphix/ source/{ref}/enable
Disable	source "name" disable	Source	SourceDisableParameters	POST /resources/ json/delphix/ source/{ref}/ disable

Environment Operations

GUI	CLI	API Topic	Input Object	Web Services
Add environment	<code>environment create</code>	SourceEnvironment	SourceEnvironmentCreateParameters	POST /resources/json/delphix/environment
Update	<code>environment "name" update</code>	SourceEnvironment	Environment	POST /resources/json/delphix/environment/{ref}
Delete	<code>environment "name" delete</code>	SourceEnvironment	-	DELETE /resources/json/delphix/environment/{ref}
Refresh	<code>environment "name" refresh</code>	SourceEnvironment	-	POST /resources/json/delphix/environment/{ref}/refresh
Enable	<code>environment "name" enable</code>	SourceEnvironment	-	POST /resources/json/delphix/environment/{ref}/enable
Disable	<code>environment "name" disable</code>	SourceEnvironment	-	POST /resources/json/delphix/environment/{ref}/disable

GUI	CLI	API Topic	Input Object	Web Services
Add manual repository	repository create	SourceRepository	SourceRepository	POST /resources/ json/delphix/ repository
Update repository	repository "name" update	SourceRepository	SourceRepository	POST /resources/ json/delphix/ repository/{ref}
Remove manual repository	repository "name" delete	SourceRepository	-	DELETE / resources/json/ delphix/ repository/{ref}
Show host details	host "name" get	Host	-	GET /resources/ json/delphix/ host/{ref}
Add cluster node	environment oracle clusternode create	OracleClusterNode	OracleClusterNode	POST /resources/ json/delphix/ environment/ oracle/ clusternode
Update cluster node	environment oracle clusternode "name" update	OracleClusterNode	OracleClusterNode	POST /resources/ json/delphix/ environment/ oracle/ clusternode/ {ref}

GUI	CLI	API Topic	Input Object	Web Services
Remove cluster node	environment oracle clusternode "name" delete	OracleClusterNode	-	DELETE /resources/json/delphix/environment/oracle/clusternode/{ref}
Add listener	environment oracle listener create	OracleListener	OracleListener	POST /resources/json/delphix/environment/oracle/listener
Update listener	environment oracle listener "name" update	OracleListener	OracleListener	POST /resources/json/delphix/environment/oracle/listener/{ref}
Remove listener	environment oracle listener "name" delete	OracleListener	-	DELETE /resources/json/delphix/environment/oracle/listener/{ref}

CLI to Python Transition

This topic describes using the CLI to understand the Python APIs.

The [command-line interface](#) is a direct translation of the web services API to an interactive environment. This allows you to use the CLI to explore functionality with tab completion, integrated help, stronger type checking, and indication of expected types and required fields. When trying to determine how to invoke an operation through the Python API or interpret the results, it is recommended that you first learn how to do the same through the CLI, and then use the provided tools to translate that into Python calls.

Installation

The Delphix Python API is available through PyPI and you may install it with pip.

```
pip install delphixpy
```

Minimum Python version

Requiring Version 2.7 and above.

Connecting to the Delphix Engine

In the Delphix Python API, all operations take an engine object which represents your connection to a Delphix Engine. Here is how you connect to the Delphix Engine using the Python API and acquire the engine object.

```
from delphixpy.delphix_engine import DelphixEngine
engine = DelphixEngine("delphix-address", "delphix-user", "delphix-password",
"DOMAIN") # Instead of DOMAIN, use SYSTEM if you are using the sysadmin user.
```

CLI Translation to Python

For backward compatibility purposes, delphixpy provides the ability to write integrations against a specific API version. The latest version is always in the root of the package. Writing against the latest version requires you to update your integrations if the API changes in future versions of the API.

Specific API versions can be used by importing the corresponding sub-package. The sub-packages are named after the API versions in the format `v<major>_<minor>_<micro>`. For example, if you want to look into API 1.5.0, you should be using modules from `delphixpy.v1_5_0`. Modules from different sub-package versions cannot interact with each other so be careful if you wish to mix API versions in the same code base.

All CLI namespaces have a corresponding Python package in which operations can be accessed. The main Python package is called `web`. All value objects which can be manipulated or read through the CLI can be found in `web.vo`.

```
delphix> database provision
```

Is equivalent to:

```
from delphixpy.web import database
database.provision(engine, provision_parameters)
```

The `provision_parameters` object in this example is an instance of `ProvisionParameters` which can be found in `delphixpy.web.vo`. The properties of the object map to the parameters you would need to specify before doing a commit in the CLI provision context.

This is an example of an "operation", as they are invoked on an object. Operations that are invoked on a particular object take a reference to that object.

```
delphix> database "dexample" refresh
```

Is equivalent to (connection code omitted):

```
from delphixpy.web import database
database.refresh(engine, "ORACLE_DB_CONTAINER-3", RefreshParameters)
```

While the CLI uses names to refer to objects, the Python API, just like the web services, use references (`ORACLE_DB_CONTAINER-3`). Persistent objects (those with a permanent unique identity) have a `reference` field that is used in all cases to refer to the object. This allows references to remain constant even if objects are renamed.

For the standard CRUD (create, read, update, delete) operations, the mapping is slightly different:

CLI Operation	Python API
<code>group list</code>	<code>group.get_all(engine)</code>
<code>group create</code>	<code>group.create(engine, group=<delphixpy.web.vo.Group>)</code>
<code>group "dexample" get</code>	<code>group.get(engine, <reference>)</code>
<code>group "dexample" update</code>	<code>group.update(engine, <reference>, group=<delphixpy.web.vo.Group>)</code>
<code>group "dexample" delete</code>	<code>group.delete(engine, <reference>)</code>

Example: Creating a Group

This is how you can create a group as a fully working example.

```
from delphixpy.web import group
```

```
from delphixpy.web.vo import Group
from delphixpy.delphix_engine import DelphixEngine
engine = DelphixEngine("delphix-address", "delphix-user", "delphix-password",
"DOMAIN")
my_group = Group()
my_group.name = "My Group"
my_group.description = "This is my new group!"
group.create(engine, my_group)
```

Asynchronous Mode

The Python API runs by default in synchronous mode. If you would wish to perform operations asynchronously there is a context manager that allows you to do that. If you need to track job progress in asynchronous mode, you can get the reference of the last job started from `engine.last_job`. When exiting the async context manager, it will wait for all jobs started within the context to finish. You can also clear the job from the context so that you do not wait for its completion or status when exiting the context manager. If a job fails, `exceptions.JobError` will be thrown.

Here is how you would perform a sync operation on all databases asynchronously.

```
from delphixpy.delphix_engine import DelphixEngine
from delphixpy import job_context
from delphixpy.web import database
engine = DelphixEngine("delphix-address", "delphix-user", "delphix-password",
"DOMAIN")
all_databases = database.get_all(engine)
with job_context.asyncly(engine):
    for db in all_databases:
        database.sync(engine, db.reference)
```

Python Cookbook: Adding a UNIX Host

This topic describes the process of adding a UNIX host using the delphixpy Python API.

Within Delphix, there are both hosts and host environments. A host represents a remote system, but may or may not be a source or target for linking or provisioning. For example, in an Oracle RAC cluster, the cluster environment represents the location of the Oracle installation(s), and while there are hosts within that cluster they are not individually manageable as environments.

Procedure

1. Create new environment creation parameters and initialize the structure for a UNIX host.

```
from delphixpy.web.vo import HostEnvironmentCreateParameters,
    UnixHostEnvironment, UnixHostCreateParameters, UnixHost, EnvironmentUser,
    PasswordCredential

host_environment_create_parameters_vo = HostEnvironmentCreateParameters()
host_environment_create_parameters_vo.host_environment = UnixHostEnvironment()
host_environment_create_parameters_vo.host_parameters =
    UnixHostCreateParameters()
host_environment_create_parameters_vo.host_parameters.host = UnixHost()
host_environment_create_parameters_vo.primary_user = EnvironmentUser()
host_environment_create_parameters_vo.primary_user.credential =
    PasswordCredential()
```

2. Set the host address and port.

The name of the environment is derived from the address used, though you can provide a more descriptive name if desired. The address can be a DNS names, IP addresses, or a comma separated list of the above.

```
host_environment_create_parameters_vo.host_parameters.host.addresses = ["192.16
8.1.2"]
host_environment_create_parameters_vo.host_parameters.host.port = 22
```

3. Set the toolkit path.

This is where Delphix will store temporary binaries used while the host is configured as part of Delphix.

```
host_environment_create_parameters_vo.host_parameters.host.toolkit_path = "/
var/delphix"
```

4. Set the username and password to use when connecting over SSH.

This user must have the privileges described in the Delphix Administration Guide. To configure a SSH user, change the credential object to `SystemKeyCredential`.

```
host_environment_create_parameters_vo.primary_user.name = "oracle"
host_environment_create_parameters_vo.primary_user.credential.password = "my
secret password"
```

5. Commit the result. A reference to your new environment will be returned from the create call. The environment discovery process will execute as an asynchronous job. The default behavior is to wait for the result, so progress will be updated until the discovery process is complete or fails.

```

from delphixpy.delphix_server import DelphixServer
from delphixpy.web import environment
server = DelphixServer("delphix-address", "delphix-user", "delphix-password",
"DOMAIN")
new_environment_reference = environment.create(server,
host_environment_create_parameters_vo)

```

6. Full example

```

from delphixpy.delphix_server import DelphixServer
from delphixpy.web import environment
from delphixpy.web.vo import HostEnvironmentCreateParameters,
UnixHostEnvironment, UnixHostCreateParameters, UnixHost, EnvironmentUser,
PasswordCredential

host_environment_create_parameters_vo = HostEnvironmentCreateParameters()
host_environment_create_parameters_vo.host_environment = UnixHostEnvironment()
host_environment_create_parameters_vo.host_parameters =
    UnixHostCreateParameters()
host_environment_create_parameters_vo.host_parameters.host = UnixHost()
host_environment_create_parameters_vo.primary_user = EnvironmentUser()
host_environment_create_parameters_vo.primary_user.credential =
    PasswordCredential()

host_environment_create_parameters_vo.host_parameters.host.addresses = ["192.16
8.1.2"]
host_environment_create_parameters_vo.host_parameters.host.port = 22
host_environment_create_parameters_vo.host_parameters.host.toolkit_path = "/
var/delphix"
host_environment_create_parameters_vo.primary_user.name = "oracle"
host_environment_create_parameters_vo.primary_user.credential.password = "my
secret password"

server = DelphixServer("delphix-address", "delphix-user", "delphix-password",
"DOMAIN")
new_environment_reference = environment.create(server,
host_environment_create_parameters_vo)

```

So You Want to Work with Delphix APIs?

What is RESTful? API? JSON? CLI? Object Reference? GET/POST? Cookies? HTTP/HTTPS? cURL? Where do I begin? Who can help me? Documentation? Tutorials? That's great for Linux, but I am on Windows? Parse? Don't have a clue about sed, awk, grep, cut, and other acronyms. What about Regular Expressions, like "Hello"?

Just a small sampling of questions that you may or may not know the answers to, let alone learning Delphix and a programming language. Delphix is a technical product, and being new to the Delphix family can be a bit overwhelming. The goal for this document is simple: to enable users to get up to speed quickly on how to use Delphix APIs.

This section covers the following topics:

- [Background Information](#)
- [Delphix API Reference URLs](#)
- [API Prerequisite Knowledge](#)
- [Delphix RESTful APIs Command Line Basics](#)
- [API Shell Scripts Programming Language Examples](#)
- [JSON Parsing](#)
- [API Use Case Commands and Scripts](#)
- [API Programming Language Examples](#)
- [API Timeflows](#)

Background Information

This document assumes that you have some basic Delphix product experience and entry-level programming knowledge. The first two sections of this document, [Delphix API Reference URLs](#) and [API Prerequisite Knowledge](#), are focused on providing the required information and reference material/URLs.

This document is for informational and demonstration purposes only. The examples are for demonstration purposes only and must be used at your own risk. As always, test and verify on development systems prior to migrating code to production environments.

What is RESTful Programming?

<http://stackoverflow.com/questions/671118/what-exactly-is-restful-programming>

- ✔ A great way to learn how to generate the Delphix RESTful API calls and the required JSON content is to use the Delphix CLI (Command Line Interface) and turn on the trace option.

```
Delphix5110HWv8> setopt trace=true
```

All subsequent CLI commands will display the GET or POST API URL with the respective input or output JSON data string. This guide will walk you through an example later.

Delphix API Reference URLs

There are a number of sources available to provide details, examples, and techniques for working with Delphix APIs. This section contains a small list of URLs that are worth reviewing/reading as required.

CLI (Command Line Interface)

The Delphix Engine provides a native command-line interface (CLI) accessible over SSH. This CLI provides an interactive layer on top of the public web service APIs, and is intended for users that wish to automate interactions with the Delphix Engine, or simply prefer a text-based interface. All of the functionality available in the CLI is also available through the public stable web service APIs should more full-featured automation be required. For more information on automation using CLI commands see [Command Line Interface Guide](#).

RESTFul APIs

Must be logged in


Users must be successfully logged in before /API pages can be accessed.

For more information on automation using the web service APIs, see the [Web Services API Guide](#).

API Documentation is also included within the Delphix Engine using the following formula:

`http://<delphix_engine>/api/`

For Example `http://172.16.160.195/api/`



For a complete list of Delphix APIs - JSON schema format, use the following URL:

http://<delphix_engine>/api/json/delphix.json

```

{
  "/delphix-about.json": {
    "cliVisibility": [
      "DOMAIN",
      "SYSTEM"
    ],
    "description": "Retrieve static system-wide properties.",
    "extends": {
      "$ref": "/delphix-typed-object.json"
    },
    "name": "PublicSystemInfo",
    "properties": {
      "apiVersion": {
        "$ref": "/delphix-version.json",
        "description": "Maximum supported API version of the current system software.",
        "type": "object"
      },
      "banner": {
        "description": "Security banner to display prior to login.",
        "type": "string"
      },
      "buildTimestamp": {
        "description": "Time at which the current system software was built.",
        "format": "date",
        "type": "string"
      },
      "buildTitle": {
        "description": "Description of the current system software.",
        "type": "string"
      },
      "buildVersion": {
        "$ref": "/delphix-version-info.json",
        "description": "Delphix version of the current system software.",
        "type": "object"
      },
      "configured": {
        "description": "Indicates whether the server has gone through initial setup or not.",
        "type": "boolean"
      },
      "currentLocale": {
        "description": "The current system locale.",
        "format": "locale",
        "type": "string"
      },
      "enabledFeatures": {
        "description": "The list of enabled features on this Delphix Engine.",
        "items": {
          "type": "string"
        },
        "type": "array"
      }
    }
  }
}

```

So, looking at the first JSON key/name:

```

"/delphix-about.json": {
  "cliVisibility": [
    "DOMAIN",
    "SYSTEM"
  ],
  "description": "Retrieve static system-wide properties.",
  . . .
}

```

And after logging into the Delphix Engine, translating this into the URL API for *about*:

<http://172.16.160.195/resources/json/delphix/about> will respond with the returned JSON data string.

```

{"type":"OKResult","status":"OK","result":{"type":"PublicSystemInfo","productType":"standard","productName":"Delphix Engine","buildTitle":"Delphix Engine 5.1.2.0","buildTimestamp":"2016-09-02T22:28:43.000Z","buildVersion":{"type":"VersionInfo","major":5,"minor":1,"micro":2,"patch":0},"configured":true,"enabledFeatures":["XPP","JETSTREAM"],"apiVersion":{"type":"APIVersion","major":1,"minor":8,"micro":0},"banner":null,"locales":["en-US"],"currentLocale":"en-US"},"job":null,"action":null}

```

For now, just remember that the Delphix Engine contains the API Documentation and Delphix JSON schema.

Masking APIs

Please refer to the Masking documentation at <https://maskingdocs.delphix.com> for information on the Masking APIs.

Cookbook Examples

Delphix documentation includes a number of cookbook examples that will not be duplicated in this section but may be referenced.

[API Cookbook: Common Tasks, Workflows, and Examples](#)

There are also working examples provided within this document and are available for download.

API Prerequisite Knowledge

JSON

JSON (JavaScript Object Notation) is a minimal, readable format for structuring data. It is a simple format for transmitting data between applications, as an alternative to XML. The Delphix API uses JSON data structure in the format of strings to send and receive data from the API calls, as you will see later in the examples. First, let's look at the JSON fundamentals.

Keys and Values

The two primary parts that makeup JSON are keys and values. Together they make key/value pairs, also called name/value pairs.

- **Key** – Always a string enclosed in quotation marks.
- **Value** – Can be a string, number, boolean expression, array, or object.
- **Key/Value Pair** – Follows a specific syntax, with the key followed by a colon followed by the value. Key/value pairs are comma separated.

Let's take a JSON sample string and identify each part of the code.

```
{
  "foo" : "bar",
  "rows" : 100
}
```

The curly brackets start and end the string. The key is "foo" and the value is "bar". A colon (:) is the delimiter between them. A comma (,) is the delimiter for multiple key/value pairs. The second pair is "rows" and the value is a number of 100.

Types of Values

Number	An integer or a decimal number
Boolean	True or false
String	Plain text alphanumeric readable characters
Null	Empty
Array	An associative array of values
Object	An associative array of key/value pairs

Numbers, Booleans, and Strings.

It is very important to understand the APIs JSON object definitions. Quoted values are treated as strings!

"x" : "1" is treated as a string, while

"x" : 1 is treated as a number

"y" : "true" is treated as a string, while

"y" : true is treated as a boolean true (false)

Null values

```
{
  "z" :
  , "b" : "World"
}
```

Nulls are empty values, but sometimes programmers code "" as a null value.

```
{
  "z" : ""
  , "b" : "World"
}
```

So always verify how the null values are defined and handled by the application.

Arrays

An array is indicated with the square brackets: [value1, value2, etc.]. In this example, we have added a categories key with an array of values.

```
...
"foo" : {
  "bar" : "Hello",
  "category" : [ "greetings", "morals" ]
}
...
```

Objects

An object is indicated by curly brackets: {"key", "value"}. Everything inside of the curly brackets is part of the object. We already learned that a value could be an object. Therefore, "foo" and the corresponding object are a key/value pair.

```
...
"foo" : { "bar" : "Hello" }
...
```

The key/value pair "bar" : "Hello" is nested inside the key/value pair "foo" : { ... }. That is an example of a hierarchy (or nested data) within JSON data.

Arrays and Objects can be nested or contained within the same level.

Summary

JSON arrays are [, ,]

JSON nested objects are , , "x":{ "a":"1", "b":"2" }, ,

JSON data can be passed within the HTTP URL (file or argument), the header, or other handlers.

From within Shell Scripts or Programming Languages, JSON data is typically processed through a "JSON parser." This topic is covered later.

Delphix CLI

Connecting to the Delphix Engine CLI

Reference: [Connecting to the CLI](#)

There are two user roles accessible, the **sysadmin** and the **delphix_admin**.

From a shell environment, you can connect using the ssh command. The IP Address (or Hostname) represents the Delphix Engine (case sensitive):

```
ssh sysadmin@127.16.160.195
```

```
ssh delphix_admin@127.16.160.195
```

From a putty session, open an ssh connection to the Delphix Engine IP Address or Hostname (case sensitive):

```
open 127.16.160.195
```

```
Login User: sysadmin@SYSTEM
```

```
#... or ...
```

```
Login User: delphix_admin@DOMAIN
```

After entering the correct password for the respective user, the menus for that user's role will be different. For example, the **sysadmin@SYSTEM** user has engine storage, whereas the **delphix_admin@DOMAIN** user has database provisioning.

You can use the CLI for scripting and configure the connection for ssh passwordless connections.

[CLI Cookbook: Configuring Key-Based SSH Authentication for Automation](#)

How to use the CLI to learn the APIs

As stated earlier, a great way to learn how to generate the Delphix RESTful API calls and the required JSON content is to use the Delphix CLI (Command Line Interface) and turn on the `CLI> setopt trace=true` option.

Below is an example of how to get the JSON required parameters for a database refresh per the type of refresh performed.

⚠ Other types or options may require other JSON parameters, so after changing any parameter, we recommend performing an "ls" command to see if there are any new parameters and/or required values.

The refresh database example below shows how to use the CLI to identify reference objects for other CLI commands and the respective RESTful API structure when the **setopt trace=true** option is set.

```

$ ssh delphix_admin@172.16.160.195
Password:
Delphix5030HWv8> ls
Children
about
action
...
connectivity
database
environment
...
toolkit
user

Operations
version
Delphix5030HWv8> database
Delphix5030HWv8 database> ls
Objects
NAME                PROVISIONCONTAINER  DESCRIPTION
DPXDEV01            -
Vdelphix_demo       delphix_demo        -
delphix_demo        -
Scripts             -
V_2C1               Scripts              -
Vvfiles             -

Children
template

Operations
createEmpty
createRestorationDataset
export
fileMapping
link
oracleSupportedCharacterSets
provision
validateXpp
xpp

```

First, we need to identify the target Delphix virtualized database object to refresh ...

 Each Delphix object has a reference that is typically used for parameter values.

```
Delphix5030HWv8 database> select Vdelphix_demo
Delphix5030HWv8 database 'Vdelphix_demo'> ls
Properties
  type: MSSqlDatabaseContainer
  name: Vdelphix_demo
  creationTime: 2016-06-16T14:30:03.033Z
  currentTimeflow: 'DB_PROVISION@2016-06-16T10:30:08'
  delphixManaged: true
  description: (unset)
  group: Windows
  masked: false
  os: Windows
  performanceMode: DISABLED
  processor: x86
  provisionContainer: delphix_demo
  reference: MSSQL_DB_CONTAINER-39
  restoration: false
  runtime:
    type: MSSqlDBContainerRuntime
    logSyncActive: false
  sourcingPolicy:
    type: SourcingPolicy
    loadFromBackup: false
    logsyncEnabled: false
  transformation: false

Operations
delete
...
purgeLogs
refresh
removeLiveSource
...
Delphix5030HWv8 database 'Vdelphix_demo'> refresh
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> ls
Properties
  type: RefreshParameters
  timeflowPointParameters:
    type: TimeflowPointSemantic
    container: (required)
    location: LATEST_POINT
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> set
timeflowPointParameters.container=delphix_demo
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> ls
Properties
  type: RefreshParameters
  timeflowPointParameters:
    type: TimeflowPointSemantic
    container: delphix_demo (*)
```

```

    location: LATEST_POINT
Delphix5030HWv8 database 'Vdelphix_demo' refresh > *commit
  Dispatched job JOB-100
  DB_REFRESH job started for "Windows/Vdelphix_demo".
  Validating that this dataset is managed by Delphix.
  Stopping virtual database.
  Unmounting datasets.
  Unexporting storage containers.
  Metadata for dSource "Vdelphix_demo" successfully deleted.
  Starting provisioning of virtual database "Vdelphix_demo".
  Creating new TimeFlow.
  Generating recovery scripts.
  Mounting datasets.
  Mounting read-only source logs dataset.
  Running user-specified pre-provisioning script.
  Recovering virtual database.
  The virtual database recovery was successful.
  Unmounting read-only source logs dataset.
  Running user-specified post-provisioning script.
  The virtual database "Vdelphix_demo" was successfully provisioned.
  DB_REFRESH job for "Windows/Vdelphix_demo" completed successfully.

```

Refresh again but this time turn on the `setopt trace=true` option.

```

Delphix5030HWv8 database 'Vdelphix_demo'> refresh
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> ls
Properties
  type: RefreshParameters
  timeflowPointParameters:
    type: TimeflowPointSemantic
    container: (required)
    location: LATEST_POINT
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> set
timeflowPointParameters.container=delphix_demo
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> ls
Properties
  type: RefreshParameters
  timeflowPointParameters:
    type: TimeflowPointSemantic
    container: delphix_demo
    location: LATEST_POINT
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> setopt trace=true
Delphix5030HWv8 database 'Vdelphix_demo' refresh *> commit
=== POST /resources/json/delphix/database/MSSQL_DB_CONTAINER-39/refresh ===
{
  "type": "RefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "MSSQL_DB_CONTAINER-38"
  }
}

```

```

    }
  }
  ...

```

⚠ The "container" value in the JSON output above is different from the target VDB reference because we are refreshing from the source database container! In this example, the **set** `timeflowPointParameters.container=delphix_demo` is represented in JSON output as **"container": "MSSQL_DB_CONTAINER-38"**

Using the CLI, you can identify the RESTful API POST and GET commands along with the JSON input data requirements.

```

=== POST /resources/json/delphix/database/MSSQL_DB_CONTAINER-39/refresh ===
{
  "type": "RefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "MSSQL_DB_CONTAINER-38"
  }
}

```

So framing the RESTful URL for a virtual database refresh, the URL will look like

http://<delphix_engine>/resources/json/delphix/database/ MSSQL_DB_CONTAINER-39 /refresh

where the **MSSQL_DB_CONTAINER-39** represents the target virtualized database to refresh. We need to POST the JSON data to the URL for processing.

```

{
  "type": "RefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "MSSQL_DB_CONTAINER-38"
  }
}

```

The **"timeflowPointParameters"** key has 6 **"type": "..."** options, each of which has its own set of parameters. The type **"TimeflowPointSemantic"** uses the default LATEST_POINT within the source container, so for simplicity, we will use this type. For more information on timeflowPointParameters 6 types, see the Advanced Section.

If this is a little confusing at this point, do not worry, that's typical. Complete examples will be shown later. The important items to remember are:

- Delphix often uses object reference names within the JSON data.

- Using the `setopt trace=true` the option provides the construct for the RESTful API URLs and the JSON data for POST / GET operations.

HTTP

We use the HTTP protocol every day for web browsing and commercial business. From finding a new restaurant to buying a 1986 Ford Thunderbird Turbo Coupe!

Most people see the HTTP within the URL Address field within the Web Browser window – for example, <http://www.google.com>

But behind the scenes, HTTP is performing a wide range of functionality. For RESTful APIs, they use HTTP's GET and POST form functionality to process data. In Delphix's case, the data is also represented as JSON structures.

HTTP GET operation is used to return data only, while HTTP POST operation is used to provide data input in the form of a structured JSON data string or file.

cURL

What is cURL?

The `cURL` client command is based on a library supporting a number of web protocols, including HTTP. The "curl" command can be called from the command line, while the cURL library is commonly integrated with your favorite programming languages, such as Java, JSP, Python, Perl, PHP, .NET, and PowerShell.

Due to its widespread adoption, we will use cURL for making the Delphix RESTful API calls within this document. Some operating systems or languages support their own HTTP commands / related libraries, and you can use these instead of cURL. One alternative is the "wget" command described later.

Is cURL installed?

```
Operating System Prompt> curl --version
curl 7.19.7 (x86_64-redhat-linux-gnu) libcurl/7.19.7 NSS/3.19.1 Basic ECC zlib/1.2.3
libidn/1.18 libssh2/1.4.2
Protocols: tftp ftp telnet dict ldap ldaps http file https ftps scp sftp
Features: GSS-Negotiate IDN IPv6 Largefile NTLM SSL libz
```

Get the HTTP output from google.com

```
Operating System Prompt> curl www.google.com
```

Wget

An alternative to cURL is Wget, which is typically a native command on all Linux environments. See the Appendix for a complete comparison between Wget and cURL.

dxtoolkit2

Delphix has developed a very robust toolkit, dxtoolkit2, which utilizes the Delphix RESTful APIs. This toolkit is cross-platform. Its commands are built with the Perl programming language.

We recommend that you review the dxtoolkit2 documentation; you may find a utility that already performs your desired function. For example, the utility `dx_get_analytics` is absolutely great for dumping analytic data from the

Delphix Engine into a .csv (comma-separated value) format, which you can then easily integrate into your enterprise monitoring tools. See the sample "Analytics" use case.

Contact Delphix personnel for the latest download.

Delphix RESTful APIs Command Line Basics

Authentication

RESTful APIs require authentication. Just plugging the URL into a web browser or running an operating system cURL command will return an authentication/login required error message.



Command Line:

```
curl http://172.16.160.195/resources/json/delphix/environment
```

Response:

```
<!DOCTYPE html><html><head><title>Apache Tomcat/8.0.29 - Error report</title><style
type="text/css">H1 {font-family:Tahoma,Arial,sans-serif;color:white;background-color:#
525D76;font-size:22px;} H2 {font-family:Tahoma,Arial,sans-
serif;color:white;background-color:#525D76;font-size:16px;} H3 {font-
family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;font-size:14px;}
BODY {font-family:Tahoma,Arial,sans-serif;color:black;background-color:white;} B
{font-family:Tahoma,Arial,sans-serif;color:white;background-color:#525D76;} P {font-
family:Tahoma,Arial,sans-serif;background:white;color:black;font-size:12px;}A
{color : black;}A.name {color : black;}.line {height: 1px; background-color: #525D76;
border: none;}</style> </head><body><h1>HTTP Status 403 - Use /resources/json/
delphix/login to log in first</h1><div class="line"></div><p><b>type</b> Status
report</p><p><b>message</b> <u>Use /resources/json/delphix/login to log in first</
u></p><p><b>description</b> <u>Access to the specified resource has been forbidden.</
u></p><hr class="line"><h3>Apache Tomcat/8.0.29</h3></body></html>
```

The authentication process requires you to establish a session first.

The example within this section illustrates the session/login and subsequent API calls with cURL using cookies created when the session was established.

Session

API Version Information

When programming for compatibility, the API version number is very important. Please be aware of the differences between versions for enterprise applications. For example, if you specify a 1.11.6 version, ONLY the available calls and functionality for that version will be used, and it will only be operational on Delphix Engine versions that support that version.

⚠ **The session uses the `-c` for the cookie creation, the login uses the `-b` for the exiting cookie created by the session, and `-c` to create a new cookie, the other commands use the `-b` for using the existing cookie created by the login.**

From the Unix/Linux command line

```
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/session \
-c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 11,
    "micro": 6
  }
}
EOF
```

Returned to the command line are the results (added linefeeds for readability)

```
{
  "status": "OK",
  "result": {
    "type": "APISession",
    "version": {
      "type": "APIVersion",
      "major": 1,
      "minor": 11,
      "micro": 6
    },
    "locale": "en_US",
    "client": null
  },
  "job": null
}
```

Login

Once you have established the session, the next step is to authenticate to the server by executing the Login Request API. Unauthenticated sessions are prohibited from making any API calls other than this login request. The username can be either a system user or domain user; the backend will authenticate using the appropriate method.

⚠ **The session uses the `-c` for the cookie creation, the login uses the `-b` for the exiting cookie created by the session, and `-c` to create a new cookie, the other commands use the `-b` for using the existing cookie created by the login.**

From the Unix/Linux command line

```
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "delphix_username",
  "password": "delphix_password"
}
EOF
```

Returned to the command line are the results (added linefeeds for readability)

```
{
  "status": "OK",
  "result": "USER-2",
  "job": null,
  "action": null
}
```

Sample Delphix API call

With a successful authentication (session, login, and saved cookie), calls to the Delphix Engine can now be made to perform the desired functionality.

⚠ The session uses the `-c` for the cookie creation, the login uses the `-b` for the exiting cookie created by the session, and `-c` to create a new cookie, the other commands use the `-b` for using the existing cookie created by the login.

For starters, let's create a session, login, and get the existing environments defined within the Delphix Engine.

```
curl -s -X POST -k --data @- http://172.16.160.195/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 7,
    "micro": 0
  }
}
EOF
```

```
curl -s -X POST -k --data @- http://172.16.160.195/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
```

```

    "username": "admin",
    "password": "delphix"
  }
EOF

```

```

curl -X GET -k http://172.16.160.195/resources/json/delphix/environment \
  -b ~/cookies.txt -H "Content-Type: application/json"

```

Returned to the command line are the results (added linefeeds for readability)

```

{
  "type": "ListResult",
  "status": "OK",
  "result":
  [
    { "type": "WindowsHostEnvironment",
      "reference": "WINDOWS_HOST_ENVIRONMENT1",
      "namespace": null,
      "name": "Window Target",
      "description": "",
      "primaryUser": "HOST_USER-1",
      "enabled": false,
      "host": "WINDOWS_HOST1",
      "proxy": null
    },
    {
      "type": "UnixHostEnvironment",
      "reference": "UNIX_HOST_ENVIRONMENT-3",
      "namespace": null,
      "name": "Oracle Target",
      "description": "",
      "primaryUser": "HOST_USER-3",
      "enabled": true,
      "host": "UNIX_HOST-3", "aseHostEnvironmentParameters": null
    }
  ],
  "job": null,
  "action": null,
  "total": 2,
  "overflow": false
}

```

Windows PowerShell Authentication Example

See the PowerShell section below if cURL is not yet available on your operating system.

These commands work on Windows Command Prompt with the respective JSON files: session.json and login.json

Filename: session.json

```
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 7,
    "micro": 0
  }
}
```

Filename: login.json

```
{
  "type": "LoginRequest",
  "username": "admin",
  "password": "delphix"
}
```

(In Powershell you use curl.exe or modify the default alias)...

```
curl.exe --insecure -c cookies.txt -sX POST -H "Content-Type: application/json" -d
"@session.json" http://172.16.160.195/resources/json/delphix/session
curl.exe --insecure -b cookies.txt -c cookies.txt -sX POST -H "Content-Type:
application/json" -d "@login.json" http://172.16.160.195/resources/json/delphix/login
curl.exe --insecure -b cookies.txt -sX GET -H "Content-Type: application/json" -k
http://172.16.160.195/resources/json/delphix/system
```

Putting the above commands within a Powershell script:

Filename: auth1.ps1

```
<#
Filename: auth.ps1
Description: Delphix Powershell Sample Authentication Script ...
Date: 2016-08-02
Author: Bitt...
#>
```

Variables ...

```
$nl = [Environment]::NewLine
$BaseURL = "http://172.16.160.195/resources/json/delphix"
$cookie = "cookies.txt"
```

Session JSON Data ...

```
write-output "${nl}Creating session.json file ..."
$json = @"
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 11,
    "micro": 6
  }
}
"@
```

Output File using UTF8 encoding...

```
write-output $json | Out-File "session.json" -encoding utf8
```

Delphix cURL Session API ...

```
write-output "${nl}Calling Session API ...${nl}"
$results = (curl.exe --insecure -c "$cookie" -sX POST -H "Content-Type: application/
json" -d "@session.json" -k $BaseURL/session)
write-output "Session API Results: ${results}"
```

Login JSON Data ...

```
write-output "${nl}Creating login.json file ..."
$user = "admin"
$pass = "delphix"
$json = @"
{
  "type": "LoginRequest",
  "username": "${user}",
  "password": "${pass}"
}
"@
```

Output File using UTF8 encoding...

```
write-output $json | Out-File "login.json" -encoding utf8
```

Delphix cURL Login API ...

```
write-output "${nl}Calling Login API ...${nl}"
$results = (curl.exe --insecure -b "$cookie" -c "$cookie" -sX POST -H "Content-Type:
application/json" -d "@login.json" -k $BaseURL/login)
write-output "Login API Results: ${results}"
```

Delphix cURL system API ...

```
write-output "${nl}Calling System API ...${nl}"
$results = (curl.exe --insecure -b "$cookie" -sX GET -H "Content-Type: application/
json" -k $BaseURL/system)
write-output "System API Results: ${results}"
```

The end is near ...

```
echo "${nl}Done ...${nl}"
exit;
```

Sample Powershell Script Output:

```
PS> . .\auth.ps1
Creating session.json file ...
Calling Session API ...
Session API Results:
{"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"API
Version","major":1,"minor":7,"micro":0},"locale":null,"client":null},"job":null,
"action":null}
Creating login.json file ...
Calling Login API ...
Login API Results: {"type":"OKResult","status":"OK","result":"USER2","job":null,"acti
on":null}
Calling System API ...
System API Results: {"type":"OKResult","status":"OK","result":{"type":"SystemInfo","p
roductType":"standard","productName":"Delphix Engine","buildTitle":"Delphix Engine
5.1.1.0","buildTimestamp":"20160721T07:23:41.000Z","buildVersion":{"type":"VersionInf
o","major":5,"minor":1,"micro":1,"patch":0},"configured":true,"enabedFeatures":["XPP",
"MSQLHOOKS"],"apiVersion":{"type":"APIVersion","major":1,"minor":8,"micro":0},"banner
":null,"locals":["enUS"],"currentLocale":"enUS","hostname":"Delphix5110HWv8","sshPubl
icKey":"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQD0srp7Aj6hFQh9yBq7273B+qtPKmCu1B18nPvr08yjt/
IZeM4qKk7caxExQS9rpfU8AWoT7e8ESV7NkBmUzOhrHnLsuJtxPqeYoqeMubVxYjJuxlH368sZuYsnB04KM0m
i39e15lxVGvxQk9tyMpl7gs7cXRz1k6puncyicU/
axGq7ALHU2uyQoVmlPasuHJbq23d21VAYLuscbtgpZLAF1R8eQH5Xqaa0RT+aQJ6B1ihZ7S0ZN914M2gZHNY
cSGDWZHwUnBGttnxx1ofRcyN4/qwT5ihq5kjApjSaNgSAU0ExqDHiqgTq0wttf5nltCqGMTFR7XY38HiNq+
+atDroot@Delphix5110HWv8\n","memorySize":8.58107904E9,"platform":"VMware with BIOS
date 05/20/2014","uuid":"564d7e1df4cb-f91098fd348d74817683","processors":[{"type":"CP
UInfo","speed":2.5E9,"cores":1}], "storageUsed":2.158171648E9,"storageTotal":2
.0673724416E10,"installationTime":"2016-07-27T13:28:46.000Z"},"job":null,"action":
null}
Done ...
PS>
```

API Shell Scripts Programming Language Examples

Why Use Shell Scripts?

Shell scripts are great tools for rapid development and validation for simple (smaller) requirements. Most development is done iteratively, and shell scripts provide immediate feedback on logic and code.

Company-supported programming languages are the preferred tools for enterprise applications.

Most likely, your company employs more Java and/or PHP programmers than Linux Shell script programmers.

The "use cases" will be programmed using either Unix/Linux Shell and/or Windows PowerShell scripts. You can easily port the logic from these scripts into your favorite programming language. Basic examples of connection with API will be provided for a number of major programming languages in a later section of this document.

Linux/Unix/(and Mac too) Shell Scripts

There are numerous shell environments, including sh, bash, csh, ksh, and tsh. Identify the current shell environment using any of the commands below:

```
ps -p $$
ps -p $$ -ocomm=
echo $0
ps -ef | grep $$ | grep -v grep
ps -ef | egrep "^\s*\d+\s+$$\s+"
```

The examples provided have all been run from the bash shell environment and may or may not run the same as the other shells. We recommend that you start a bash shell by typing **bash** at the operating system prompt, like this:

```
Operating_System_Prompt> bash
```

The scripts included within this document have all been run on Linux and Mac environments within a bash shell and are NOT certified by any means. As always, test and verify in development for your environment.

For non-Linux platforms and non-bash shell environments, please re-validate for your configuration and search the web for any alternative methods/tools/utilities that may perform the same actions.

Windows PowerShell

Powershell Open Source is now available for Linux and Mac OS. for more information refer to [Microsoft Open Source](#).

Requirements

Windows has a number of versions of Powershell. The minimum version for Delphix is 2.0 for SQL Server 2008 environments. There are numerous enhancements and features with subsequent Powershell versions. Additionally, you must be aware of the architecture of 32bit or 64bit Powershell versions you are running from within.

```
PS> $PSVersionTable.PSVersion
```

Major	Minor	Build	Revision
2	0	-1	-1

32bit or 64 bit

If executing Powershell scripts from within Delphix Pre/Post Scripts commands or Delphix hooks, the default Powershell used is 32 bit, whereas the typical default Windows Powershell is 64 bit. However, Powershell allows you to execute 64 bit Powershell command from within the 32 bit environment. Shown below is a simple alias, ps64, to execute 64bit Powershell scripts.

```
PS> set-alias ps64 "$env:windir\sysnative\WindowsPowerShell\v1.0\powershell.exe"
```

Sample call to execute 64bit Powershell script

```
PS> ps64 [path\to\any_64bit_powershell_script].ps1
```

Courtesy of this article: <http://www.gregorystrike.com/2011/01/27/how-to-tell-if-powershell-is-32-bit-or-64-bit/>

```
PS> if ($env:Processor_Architecture -eq "x86") { write "running on 32bit" } else { write "
running on 64bit" }
running on 32bit
```

...Or...

```
PS> if ([System.IntPtr]::Size -eq 4) { "32-bit" } else { "64-bit" }
32-bit
```

It is worth noting that the locations of the 32-bit and 64-bit versions of Powershell are somewhat misleading. The 32-bit PowerShell is found at C:\Windows\SysWOW64\WindowsPowerShell\v1.0\powershell.exe and the 64-bit PowerShell is at C:\Windows\System32\WindowsPowerShell\v1.0\powershell.exe

Execution of Scripts Security Disabled

It is possible to disable Powershell environments on the system. If they are disabled, you will see the following error for any Powershell script that you try to execute.

```

PS> . .
[any_powershell_script].ps1
File [any_powershell_script].ps1 cannot be loaded because the execution of scripts is
disabled on this system. Please see "get-help about signing" for more details.
At line:1 char:2
+ . <<<< .\ [any_powershell_script].ps1
    + CategoryInfo          : NotSpecified: ( [], PSSecurityException
    + FullyQualifiedErrorId : RuntimeException

```

To enable Powershell scripts to be executed, set the execution policy to Yes.

```

PS> set-executionpolicy remotesigned
Execution Policy Change
The execution policy helps protect you from scripts that you do not trust. Changing the
execution policy might expose you to the security risks described in the
about_Execution_Policies help topic. Do you want to change the execution policy?
[Y] Yes      [N] No      [S] Suspend    [?] Help (default is "Y"): Y
PS>

```

Now your Powershell scripts will be executed.

curl.exe

Not all Windows platforms have the cURL executable installed, the easiest way to install and use cURL on Windows is as follows:

1. Download the package from <https://curl.se/windows/> and unzip.
2. In the "bin" folder find "curl.exe" and move it to "C:\Windows\System32".

Then you will be able to use the curl command from the Windows Command Prompt or PowerShell console.

```

PS> curl.exe --version
curl 7.49.1 (x86_64-w64-mingw32) libcurl/7.49.1 OpenSSL/1.0.2h zlib/1.2.8 libidn/1.32
libssh2/1.7.0 nghttp2/1.12.0 librtmp/2.3
Protocols: dict file ftp ftps gopher http https imap imaps ldap ldaps pop3 pop3s rtmp
rtsp scp sftp smtp smtps telnet tftp
Features: IDN IPv6 Largefile SSPI Kerberos SPNEGO NTLM SSL libz TLS-SRP HTTP2 Metalink

```

Invoking the curl or curl.exe from Powershell command line.

```
PS> Get-Command curl
```

CommandType	Name	ModuleName
Alias	curl -> Invoke-WebRequest	

```
PS> Get-Command curl.exe
```

CommandType	Name	ModuleName
Application	curl.exe	

If the alias curl name is to the Invoke-WebRequest, you will need to use the curl.exe command explicitly or remove the alias.

```
PS> Remove-item alias:curl
```

Verify that curl and/or curl.exe work from the respective Powershell environment:

```
PS> curl.exe --version
curl 7.49.1 (x86_64-w64-mingw32) ...
```

```
PS> curl --version
curl 7.49.1 (x86_64-w64-mingw32) ...
```

JSON Parsing

Unix/Linux/Mac Shell

Unix/Linux tools come natively with a host of shell utilities that one can use for parsing out the desired name/value pairs. Tools include sed, awk, cut, tr, and grep, to name a few. System administrators use these utilities frequently and may be able to assist with the methods for parsing JSON strings. For more information please refer to [Parsing JSON with UNIX tools](#) and [Extract a JSON value from a BASH script](#).

Basic awk and sed parsing

```
json='{ "type": "OKResult", "status": "OK", "result":
{ "type": "Job", "reference": "JOB-53", "namespace": null, "name": null, "actionType": "DB_SYNC
", "target": "ORACLE_DB_CONTAINER-9", "targetObjectType": "OracleDatabaseContainer", "jobS
tate": "RUNNING", "startTime": "2016-08-12T19:58:59.811Z", "updateTime": "2016-08-12T19:58
:59.828Z", "suspendable": true, "cancelable": true, "queued": false, "user": "USER-2", "emailA
ddresses": null, "title": "Run SnapSync for database
\"VDPXDEV1\".", "percentComplete": 0.0, "targetName": "Oracle_Source/VDPXDEV1", "events":
[ { "type": "JobEvent", "timestamp": "2016-08-12T19:58:59.840Z", "state": null, "percentCompl
ete": 0.0, "messageCode": "event.job.started", "messageDetails": "DB_SYNC job started for
\"Oracle_Source/
VDPXDEV1\".", "messageAction": null, "messageCommandOutput": null, "diagnoses":
[ ], "eventType": "INFO" } ], "parentActionState": "WAITING", "parentAction": "ACTION-238"}, "j
ob": null, "action": null }'
```

```
echo $json | sed -e 's/[{}]/'/'/g' | awk -v RS=',,' -F: '{print $1 $2}'
"type" "OKResult"
"status" "OK"
"result" "type"
"reference" "JOB-53"
"namespace" null
"name" null
"actionType" "DB_SYNC"
"target" "ORACLE_DB_CONTAINER-9"
"targetObjectType" "OracleDatabaseContainer"
"jobState" "RUNNING"
"startTime" "2016-08-12T19
"updateTime" "2016-08-12T19
"suspendable" true
"cancelable" true
"queued" false
"user" "USER-2"
"emailAddresses" null
"title" "Run SnapSync for database \"VDPXDEV1\"."
"percentComplete" 0.0
"targetName" "Oracle_Source/VDPXDEV1"
"events" [ "type"
"timestamp" "2016-08-12T19
"state" null
"percentComplete" 0.0
```

```
"messageCode"event.job.started"
"messageDetails"DB_SYNC job started for \"Oracle_Source/VDPXDEV1\"."
"messageAction"null
"messageCommandOutput"null
"diagnoses"[]
"eventType"INFO"]
"parentActionState"WAITING"
"parentAction"ACTION-238"
"job"null
"action"null
```

Find jobState. Print the second argument, and remove the double-quotes.

```
echo $json | sed -e 's/[{}]/'/g' | sed s/\\"//g | awk -v RS=',' -F: '$1=="jobState"{p
rint $2}'
RUNNING
```

The first sed removed the brackets and braces. The second sed removes the double-quotes. The awk command parses the line by comma delimiters and then parses each line by the semi-colon delimiter and if the first variable \$1 is equal to the **jobState** value then print the second \$2 variable.

If the results contain an array of values, then you need to loop through each set and parse out the desired value. For example,

```
json='
{"type":"ListResult","status":"OK","result":[{"type":"WindowsHostEnvironment","refere
nce":"WINDOWS_HOST_ENVIRONMENT-1","namespace":null,"name":"Window Target","descriptio
n":"","primaryUser":"HOST_USER-1","enabled":false,"host":"WINDOWS_HOST-1","proxy":
null},{type":"UnixHostEnvironment","reference":"UNIX_HOST_ENVIRONMENT-3","namespace":n
ull,"name":"Oracle Target","description":"","primaryUser":"HOST_USER-3","enabled":tru
e,"host":"UNIX_HOST-3","aseHostEnvironmentParameters":null}],job":null,"action":null
,"total":2,"overflow":false}'
```

Parse out array object into separate lines

```
SOURCE_ENV="Oracle Target"
lines=`echo ${json} | cut -d "[" -f2 | cut -d "]" -f1 | awk -v RS='},{}' -F: '{print
$0}' `
while read -r line
do
    #echo "Processing $line"
    #echo $line | sed -e 's/[{}]/'/g' | sed s/\\"//g | awk -v RS=',' -F: '$1=="name"{
print $2}'
    TMPNAME=`echo $line | sed -e 's/[{}]/'/g' | sed s/\\"//g | awk -v RS=',' -F:
'$1=="name"{print $2}' `
    #echo "Name: |${TMPNAME}| |${SOURCE_ENV}|"
    if [[ "${TMPNAME}" == "${SOURCE_ENV}" ]]
    then
        echo $line | sed -e 's/[{}]/'/g' | sed s/\\"//g | awk -v RS=',' -F: '$1=="pri
maryUser"{print $2}'
```

```

        PRI_USER=`echo $line | sed -e 's/[{}]/''/g' | sed s/\\''//g | awk -v RS=',' -F:
'$1=="primaryUser"{print $2}' `
        break
    fi
done <<< "$(echo -e "$lines")"

echo "primaryUser reference: ${PRI_USER}"

```

Output:

```
primaryUser reference: HOST_USER-3
```

The above methods will be used within the sample scripts since they use the native Linux tools. They typically do not require you to load extra packages or libraries onto the system.

There are a number of open-source utilities designed to simplify the parsing of JSON, such as *jsawk* and *jq*.

jsawk

Linux:

- [jsawk](#)

Mac:

- [jsawk on MAC](#)

jq

Reference- <https://stedolan.github.io/jq/>

64-bit system:

- `wget https://github.com/stedolan/jq/releases/download/jq-1.6/jq-linux64`
- `chmod +x ./jq`
- `sudo cp jq /usr/local/bin`

Older versions:

- Reference- <https://stedolan.github.io/jq/download/>

Another method is to use an existing programming language typically available with your native operating systems, such as Perl or Python.

```

$ which perl
/usr/bin/perl

$ which python
/usr/bin/python

```

Example: Use Python to pretty format the JSON data string.

Pretty JSON using Python ...

```

json='{
  "type": "OKResult",
  "status": "OK",
  "result": {
    "type": "SystemInfo",
    "productType": "standard",
    "productName": "Delphix Engine",
    "buildTitle": "Delphix Engine 5.1.1.0",
    "buildTimestamp": "20160721T07:23:41.000Z",
    "buildVersion": {
      "type": "VersionInfo",
      "major": 5,
      "minor": 1,
      "micro": 1,
      "patch": 0,
      "configured": true,
      "enabledFeatures": ["XPP", "MSSQLHOOKS"],
      "apiVersion": {
        "type": "APIVersion",
        "major": 1,
        "minor": 8,
        "micro": 0,
        "banner": null,
        "locals": ["enUS"],
        "currentLocale": "enUS",
        "hostname": "Delphix5110HWv8",
        "sshPublicKey": "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDsrp7Aj6hFQh9yBq7273B+qtPKmCu1B18nPvr08yjt/IZeM4qKk7caxExQS9rpFU8AWoT7e8ESV7NkBMUz0HrHnLsuJtxPqeYoqeMubVxYjJuxlH368sZuYsnB04KM0mi39e15lxVGvxQk9tyMpl7gs7cXRz1k6puncyicZU/axGq7ALHU2uyQoVmlPasuHJbq23d21VAYLuscbtgpZLAF1R8eQH5Xqaa0RT+aQJ6B1ihZ7S0ZN914M2gZHNYcSGDWZHwUnBGttnxx1ofRcyN4/qwT5iHq5kjApjSaNgSAU0ExqDHiqTq0wttf5nltCqGMTFR7XY38HiNq+atDroot@Delphix5110HWv8\n",
        "memorySize": 8.58107904E9,
        "platform": "VMware with BIOS date 05/20/2014",
        "uuid": "564d7e1df4cb-f91098fd348d74817683",
        "processors": [{"type": "CPUInfo", "speed": 2.5E9, "cores": 1}],
        "storageUsed": 2.158171648E9,
        "storageTotal": 2.0673724416E10,
        "installationTime": "2016-07-27T13:28:46.000Z"},
        "job": null,
        "action": null}
  }
}'

```

Pipe the JSON data to Python programming language to pretty up the format the output for the \$json string/data.

```

$ echo $json | python -mjson.tool
{
  "action": null,
  "job": null,
  "result": {
    "apiVersion": {
      "major": 1,
      "micro": 0,
      "minor": 8,
      "type": "APIVersion"
    },
    "banner": null,
    "buildTimestamp": "20160721T07:23:41.000Z",
    "buildTitle": "Delphix Engine 5.1.1.0",
    "buildVersion": {
      "major": 5,
      "micro": 1,
      "minor": 1,
      "patch": 0,
      "type": "VersionInfo"
    },
    "configured": true,
    "currentLocale": "enUS",
    "enabledFeatures": [
      "XPP",
      "MSSQLHOOKS"
    ],
    "hostname": "Delphix5110HWv8",
    "installationTime": "2016-07-27T13:28:46.000Z",
    "locals": [

```

```

    "enUS"
  ],
  "memorySize": 8581079040.0,
  "platform": "VMware with BIOS date 05/20/2014",
  "processors": [
    {
      "cores": 1,
      "speed": 2500000000.0,
      "type": "CPUInfo"
    }
  ],
  "productName": "Delphix Engine",
  "productType": "standard",
  "sshPublicKey": "ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQD0srp7Aj6hFQh9yBq7273B+qtPKmCu1B18nPvr08yjt/
IZeM4qKk7caxExQS9rpFU8AWoT7e8ESV7NkBmUz0HrHnLsuJtxPqeYoqeMubVxYjJuxlH368sZuYsnB04KM0m
i39e15lxVGvxQk9tyMpl7gs7cXRz1k6puncyicZU/
axGq7ALHU2uyQoVmlPasuHJbq23d21VAYLuscbtgpZLAFrR8eQH5Xqaa0RT+aQJ6B1ihZ7S0ZN914M2gZHNY
cSGDWZhwUnBGttnxx1ofRcyN4/qwT5iHq5kjApjSaNgSAU0ExqDHiqgTq0wttf5nltCqGMTFR7XY38HiNq+
+atDroot@Delphix5110HWv8\n",
  "storageTotal": 20673724416.0,
  "storageUsed": 2158171648.0,
  "type": "SystemInfo",
  "uuid": "564d7e1df4cb-f91098fd348d74817683"
},
  "status": "OK",
  "type": "OKResult"
}

```

jq parser

The "jq" command line parser is available on Unix, Linux, Mac, and Windows platforms. Typically, for Windows, the built-in ConvertFrom/To-Json object parser will be used. "jq" is being included in most native Linux distributions and is easy to install on the Mac OS.

References:

[Parsing JSON with jq](#)

[Installation FAQ](#)

[Mac Installation](#)

Example:

```

json='{"type":"ListResult","status":"OK","result":
[{"type":"OracleLinkedSource","reference":"ORACLE_LINKED_SOURCE-52","namespace":null,"name":"DPXDEV0
1","description":null,"virtual":false,"restoration":false,"staging":false,"container":"ORACLE_DB_CON
TAINER-120","config":"ORACLE_SINGLE_CONFIG-40","status":"DEFAULT","runtime":
{"type":"OracleSourceRuntime","status":"RUNNING","accessible":true,"databaseSize":2.409529344E9,"not
AccessibleReason":null,"databaseMode":"READ_WRITE","lastNonLoggedLocation":"","activeInstances":
[{"type":"OracleActiveInstance","instanceNumber":1,"instanceName":"DPXDEV01","hostName":"linuxtarget
.delphix.local"}],"databaseStats":null,"bctEnabled":true,"racEnabled":null,"dnfsEnabled":false,"arch
ivelogEnabled":null},"backupLevelEnabled":false,"rmanChannels":2,"filesPerSet":5,"checkLogical":fals
e,"externalFilePath":null,"encryptedLinkingEnabled":false,"compressedLinkingEnabled":true,"bandwidth
Limit":0,"numberOfConnections":1,"enabled":true,"preScript":"","postScript":"","role":"PRIMARY"},
{"type":"OracleVirtualSource","reference":"ORACLE_VIRTUAL_SOURCE-25","namespace":null,"name":"VBITT"
,"description":null,"virtual":true,"restoration":false,"staging":false,"container":"ORACLE_DB_CONTAI
NER-121","config":"ORACLE_SINGLE_CONFIG-47","status":"DEFAULT","runtime":

```

```
{
  "type": "OracleSourceRuntime",
  "status": "RUNNING",
  "accessible": true,
  "databaseSize": 2.410053632E9,
  "notAccessibleReason": null,
  "databaseMode": "READ_WRITE",
  "lastNonLoggedLocation": "0",
  "activeInstances": [
    {
      "type": "OracleActiveInstance",
      "instanceNumber": 1,
      "instanceName": "VBITT",
      "hostName": "linuxtarget.delphix.local"
    }
  ],
  "databaseStats": [
    {
      "type": "OracleDatabaseStatsSection",
      "sectionName": "Open Transactions",
      "columnHeaders": ["Transaction Count"],
      "rowValues": [
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["0"]
        }
      ]
    },
    {
      "type": "OracleDatabaseStatsSection",
      "sectionName": "Session Statistics",
      "columnHeaders": ["Current Session", "Total Session", "High Watermark"],
      "rowValues": [
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["2", "46", "5"]
        }
      ]
    },
    {
      "type": "OracleDatabaseStatsSection",
      "sectionName": "Top Wait Events",
      "columnHeaders": ["Event", "Wait Count", "Total Wait Time (s)"],
      "rowValues": [
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["Disk file operations I/O", "13", "13"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["log file sequential read", "11", "12"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["control file parallel write", "8", "8"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["control file sequential read", "6", "3"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["ARCH wait for process start 3", "2", "2"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["db file sequential read", "9", "1"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["rdbms ipc reply", "1", "1"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["JS coord start wait", "1", "1"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["os thread startup", "2", "0"]
        },
        {
          "type": "OracleDatabaseStatistic",
          "statisticValues": ["Parameter File I/O", "1", "0"]
        }
      ]
    },
    {
      "type": "OracleDatabaseStatsSection",
      "sectionName": "Top SQL by CPU",
      "columnHeaders": ["Percentage of Load", "SQL Statement"],
      "rowValues": [
        []
      ]
    }
  ],
  "bctEnabled": false,
  "racEnabled": null,
  "dnfsEnabled": false,
  "archiveLogEnabled": null,
  "operations": [
    {
      "type": "VirtualSourceOperations",
      "configureClone": [],
      "preRefresh": [],
      "postRefresh": []
    }
  ],
  "mountBase": "/mnt/provision",
  "fileMappingRules": null,
  "manualProvisioning": null,
  "configParams": {
    "memory_target": "1191182336",
    "processes": "150",
    "log_archive_dest_1": "location=/mnt/provision/VBITT/archive/MANDATORY",
    "_omf": "ENABLED",
    "filesystemio_options": "setall",
    "compatible": "11.2.0.4.0",
    "audit_trail": "NONE",
    "remote_login_passwordfile": "EXCLUSIVE",
    "open_cursors": "300",
    "audit_sys_operations": "FALSE"
  },
  "configTemplate": null,
  "nodeListenerList": [],
  "enabled": true,
  "role": "PRIMARY"
},
{
  "job": null,
  "action": null,
  "total": 2,
  "overflow": false
}
]
```

We have a very big JSON string above. Let's perform some basic jq parsing.

1. Pipe JSON string into jq command line parser.

```
echo $json | jq '.'
```

The output is a pretty human-readable JSON formatted string.

2. Get the first-level status value (...,"status":"OK",...)

```
echo $json | jq '.status'
"OK"
```

3. Get raw values (not quoted).

```
echo $json | jq --raw-output '.status'
OK
```

4. Get a number of rows returned for the type equal to "ListResult" API returned request.

```
echo $json | jq --raw-output '.total'
2
```

5. Get the first result set.

```

echo $json | jq '.result[0] '
{
  "type": "OracleLinkedSource",
  "reference": "ORACLE_LINKED_SOURCE-52",
  "namespace": null,
  "name": "DPXDEV01",
  "description": null,
  "virtual": false,
  "restoration": false,
  "staging": false,
  "container": "ORACLE_DB_CONTAINER-120",
  "config": "ORACLE_SINGLE_CONFIG-40",
  "status": "DEFAULT",
  "runtime": {
    "type": "OracleSourceRuntime",
    "status": "RUNNING",
    "accessible": true,
    "databaseSize": 2409529344,
    "notAccessibleReason": null,
    "databaseMode": "READ_WRITE",
    "lastNonLoggedLocation": "0",
    "activeInstances": [
      {
        "type": "OracleActiveInstance",
        "instanceNumber": 1,
        "instanceName": "DPXDEV01",
        "hostName": "linuxtarget.delphix.local"
      }
    ],
    "databaseStats": null,
    "bctEnabled": true,
    "racEnabled": null,
    "dnfsEnabled": false,
    "archiveLogEnabled": null
  },
  "backupLevelEnabled": false,
  "rmanChannels": 2,
  "filesPerSet": 5,
  "checkLogical": false,
  "externalFilePath": null,
  "encryptedLinkingEnabled": false,
  "compressedLinkingEnabled": true,
  "bandwidthLimit": 0,
  "numberOfConnections": 1,
  "enabled": true,
  "preScript": "",
  "postScript": "",
  "role": "PRIMARY"
}

```

6. Get the first result set name value.

```
echo $json | jq --raw-output '.result[0].name'
DPXDEV01
```

7. Get first result set reference value.

```
echo $json | jq --raw-output '.result[0].reference'
```

8. Get first result set name=value pairs.

```
echo $json | jq '.result[0]' | jq -r "to_entries|map(\"(.key)=(.value|tostring)
\").[]" | grep container
```

```
container=ORACLE_DB_CONTAINER-120
```

9. Get ALL result sets name values.

```
echo $json | jq '.result[].name'
"DPXDEV01"
"VBITT"
```

10. Get ALL result sets "reference" and "container" values.

```
echo $json | jq '.result[].reference,.result[].container'
"ORACLE_LINKED_SOURCE-52"
"ORACLE_VIRTUAL_SOURCE-25"
"ORACLE_DB_CONTAINER-120"
"ORACLE_DB_CONTAINER-121"
```

11. Now, let's scan ALL result sets for a conditional match and return a related value.

```
echo $json | jq --raw-output '.result[] | select(.name=="VBITT") | .container'
ORACLE_DB_CONTAINER-121
```

```
echo $json | jq --raw-output '.result[] | select(.name=="VBITT") | .reference'
ORACLE_VIRTUAL_SOURCE-25
```

```
echo $json | jq --raw-output '.result[] | select(.name=="VBITT")
| .container, .reference'
ORACLE_DB_CONTAINER-121
ORACLE_VIRTUAL_SOURCE-25
```

This is the typical usage for Delphix, where the human-readable name is provided and we need to look up the object reference, container, status, etc. for the respective name. Some object references are based on expressions such as "and" or "or" conditions.

```
echo $json | jq --raw-output '.result[] |
select(.environment=="UNIX_HOST_ENVIRONMENT-9" and .name=="/u02/ora/app/product/
11.2.0/dbhome_1" ) | .reference '
```

In this case, the jq select command has an "and" condition in order to correctly identify the target result object index. This is important for getting the correct and single return value for `| .reference`, since there might be more than one instance within the environment.

For a working example of using the jq JSON parser, see the VDB Init using jq command-line JSON Parser use case, **Filename: vdb_init.sh**. A version of all the Unix/Linux/Mac shell scripts exists within the code provided. It contains the ***_jq.sh** within the filename.

PowerShell

Starting with Powershell 3.0, there are ConvertFrom-Json and ConvertTo-Json modules/commands to parse the JSON string data to/from objects. If you are stuck with Powershell 2.x., the next section provides similar functions as a method of working with JSON strings.

 These 2.x functions are not 100% the same as the Powershell 3.0 ConvertFrom-Json/ConvertTo-Json modules.

PowerShell 2 Example

Filename: parse_2.0.ps1

For Powershell 2.0, there are no JSON-provided functions or commands, so the following will serialize the JSON data to a serialized array.

```
function ConvertTo-Json20([object] $item){
    add-type -assembly system.web.extensions
    $ps_js=new-object system.web.script.serialization.javascriptSerializer
    return $ps_js.Serialize($item)
}

function ConvertFrom-Json20([object] $item){
    add-type -assembly system.web.extensions
    $ps_js=new-object system.web.script.serialization.javascriptSerializer
    # The comma operator is the array construction operator in PowerShell
    return ,$ps_js.DeserializeObject($item)
}
```

Use the JSON from the system API Call.

```
$json='{"type":"OKResult","status":"OK","result":
{"type":"SystemInfo","productType":"standard","productName":"Delphix
Engine","buildTitle":"Delphix Engine
5.1.1.0","buildTimestamp":"20160721T07:23:41.000Z","buildVersion":
{"type":"VersionInfo","major":5,"minor":1,"micro":1,"patch":0},"configured":true,"ena
```

```
bedFeatures":["XPP","MSSQLH00KS"],"apiVersion":
{"type":"APIVersion","major":1,"minor":8,"micro":0},"banner":null,"locals":
["enUS"],"currentLocale":"enUS","hostname":"Delphix5110HWv8","sshPublicKey":"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQD0srp7Aj6hFQh9yBq7273B+qtPKmCu1B18nPvr08yjt/
IZeM4qKk7caxExQS9rpFU8AWoT7e8ESV7NkBmUzOHRHnLsuJtxPqeYoqeMubVxYjJuxlH368sZuYsnB04KM0m
i39e15lxVGvxQk9tyMpl7gs7cXRz1k6puncyiczU/
axGq7ALHU2uyQoVmlPasuHJbq23d21VAYLuscbtgpZLAF1R8eQH5Xqaa0RT+aQJ6B1ihZ7S0ZN914M2gZHHNY
cSGDWZHwUnBGttnxx1ofRcyN4/qwT5iHq5kjApjSaNgSAU0ExqDHiqTq0wttf5nltCqGMTFR7XY38HiNq+
+atDroot@Delphix5110HWv8\n","memorySize":8.58107904E9,"platform":"VMware with BIOS
date 05/20/2014","uuid":"564d7e1df4cb-f91098fd348d74817683","processors":
[{"type":"CPUInfo","speed":2.5E9,"cores":1}],"storageUsed":2.158171648E9,"storageTota
l":2.0673724416E10,"installationTime":"2016-07-27T13:28:46.000Z"},"job":null,"action"
:null}'
```

Convert the JSON string.

 The job and action are null values.

```
PS> $o = ConvertFrom-Json20 $json
PS> $o

Key          Value
---          -
type         OKResult
status       OK
result       {[type, SystemInfo], [productType, standard], [productNa...
job
action
```

Extract the result JSON string array.

```
PS> $a = $o.result
PS> $a

Key          Value
---          -
type         SystemInfo
productType  standard
productName  Delphix Engine
buildTitle   Delphix Engine 5.1.1.0
buildTimestamp 20160721T07:23:41.000Z
buildVersion  {[type, VersionInfo], [major, 5], [minor, 1], [micro, 1]...
configured   True
enabledFeatures {XPP, MSSQLH00KS}
apiVersion   {[type, APIVersion], [major, 1], [minor, 8], [micro, 0]}
banner
locals       {enUS}
currentLocale enUS
hostname     Delphix5110HWv8
sshPublicKey  ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQD0srp7Aj6hFQh9yBq7...
```

```

memorySize      8581079040
platform        VMware with BIOS date 05/20/2014
uuid            564d7e1df4cb-f91098fd348d74817683
processors      {System.Collections.Generic.Dictionary`2[System.String,S...
storageUsed     2158171648
storageTotal    20673724416
installationTime 2016-07-27T13:28:46.000Z

```

Same output as above.

```

PS> foreach ($element in $a) {$element}

PS> $a.type
SystemInfo
PS> $a.buildTitle
Delphix Engine 5.1.1.0
PS> $a.hostname
Delphix5110HWv8

```

Extract the result.buildVersion object.

```

PS> $a1 = $o.result.buildVersion
PS> $a1

Key          Value
---          -
type         VersionInfo
major        5
minor        1
micro        1
patch        0

PS> $a1.major
5

```

Extract the result.processors array collection.

```

PS> $b = $o.result.processors
PS> $b

Key          Value
---          -
type         CPUInfo
speed        25000000000
cores        1

PS> $a -is [Array]
False
PS> $a -is [Object]

```

```
True
PS> $b -is [Array]
True
```

Convert Array Collection to Object.

```
PS> $b1 = $b | Select-Object
PS> $b1

Key                Value
---                -
type               CPUInfo
speed              2500000000
cores              1

PS> $b1.type
CPUInfo
PS> $b1.speed
2500000000
```

PowerShell 3 or greater Example

Starting with Powershell 3.0, there is are `ConvertFrom-Json` and `ConvertTo-Json` commands to parse the JSON data to/from objects.

Reference:

- [ConvertFrom-Json](#)
- [ConvertTo-Json](#)

```
$o = $json | ConvertFrom-Json
```

There are a number of tutorials and functional examples on the web. Below is an excerpt from the [PowerShell introduction video for Linux / Mac Open Source announcement](#).

PowerShell JSON ConvertTo-Json and Python Example 15:55 through 21:16

The concept is straightforward:

- The `ConvertFrom-Json` JSON string is converted into a Powershell object that you can reference directly.
- The `ConvertTo-Json` takes the JSON object and converts it to a string.

JSON Parsing from within Programming Languages

Most programming languages provide their own libraries, functions, and methods for parsing JSON data strings into objects/hashtes/arrays/xml that the native programming language can easily process.

API Use Case Commands and Scripts

Sample Script Parsers

The Delphix Use Cases scripts provided use the native operating system "curl" command and then a JSON parser, depending on the engine on which you are running the scripts.

For Unix/Linux/Mac, the scripts provided use both native shell commands and/or the jq parser program commands. Subroutines have been provided for both methods:

```
Native Shell commands:    parseJSON_subroutines.sh
jq Parser commands:      jqJSON_subroutines.sh
```

For Windows, the scripts are for Powershell 2.0 and utilize the custom `ConvertFrom-Json20` and `ConvertTo-Json20` functions provided. As noted, with Powershell 3.0, there are `ConvertFrom-Json` and `ConvertTo-Json` command-lets provided by Powershell.

Using the jq Parser

These are some of the jq commands used within the scripts. The first is a shell script subroutine which is used for finding and returning a single item value.

Filename: jqJSON_subroutines.sh

Subroutines

This code requires the jq Linux/Mac JSON parser program

```
jqParse() {
  STR=$1          # json string
  FND=$2          # name to find
  RESULTS=""      # returned name value
  RESULTS=`echo $STR | jq --raw-output '."'$FND'`
  #echo "Results: ${RESULTS}"
  if [ "${FND}" == "status" ] && [ "${RESULTS}" != "OK" ]
  then
    echo "Error: Invalid Satus, please check code ... ${STR}"
    exit 1;
  elif [ "${RESULTS}" == "" ]
  then
    echo "Error: No Results ${FND}, please check code ... ${STR}"
    exit 1;
  fi
  echo "${RESULTS}"
}
```

The subroutine is called from the shell script to return values based on the key (or name) value provided.

Usage – After every curl command, check that the returned status value is "OK".

```
RESULTS=$( jqParse "${STATUS}" "status" )
```

Call the jqParse subroutine, where

- \${STATUS} is the returned JSON string from the curl command, and
- the value we want returned is where the name/key is equal to "status"

Usage – Get a single value within the returned nested result object:

```
JOBSTATE=$( jqParse "${JOB_STATUS}" "result.jobState" )
```

Call the jqParse subroutine, where

- \${JOB_STATUS} is the returned JSON string from the cURL command, and
- the value we want returned is where the name/key is equal to "jobState" with the nested ".result" object.

Usage – Find name/value result object and return another value within the select result object:

Use jq to parse out container reference for name of \$SOURCE_SID ...

```
CONTAINER_REFERENCE=`echo ${STATUS} | jq --raw-output '.result[] | select(.name=="${SOURCE_SID}""") | .reference `
```

where

- \${STATUS} is the returned JSON string from the cURL command, and
- the value we want returned is based on the selected nested result object where the .result[].name is equal to "\${SOURCE_SID}" and return the .reference value for the selected result object.

Delphix Engine Use Cases

Delphix User Session Timeout

Some activities can take longer than the default 30 minute session timeout value. Therefore, the following script allows you to change the timeout value using the RESTful API. As always, you can change it easily through the CLI.

 This code is the first example showing how object references are used for input (either JSON or URL) into API calls. The name will be the **DE_USER** variable value **delphix_admin**. The object reference that the code identifies is **USER-2**, which in this case is passed into the API URL to update the user parameters passed via the JSON string.

Filename: user_timeout.sh

Edit the file to update the parameters as required for your environment.

```
#####
# DELPHIX CORP #
```

```
#####
#Parameter Initialization
DMIP=172.16.160.195
#DMPORT=8282
DMUSER=delphix_admin
DMPASS=delphix
COOKIE=~/.cookies.txt
COOKIE=`eval echo $COOKIE`
CONTENT_TYPE="Content-Type: application/json"
BaseURL="http://${DMIP}/resources/json/delphix"
#
Required for user timeout ...
#
DE_USER="delphix_admin" # Delphix Engine User
DE_TIMEOUT=120 # Timeout integer in minutes
#####
# NO CHANGES REQUIRED BELOW THIS POINT #
#####
```

Sample Output

```
$ ./user_timeout.sh # or ./user_timeout_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
user reference: USER-2
Update delphix_admin session timeout value to 120 minutes ...
Returned JSON: {"type":"OKResult","status":"OK","result":"","job":null,"action":"ACTION-423"}
Results: OK
Done ...
$
```

VDB Init (start | stop | enable | disable | status | delete)

This script is used to start, stop, enable, disable, and delete a Delphix platform source object. Typically, this is done on a virtual databases (VDBs), but you can use it for dSources as well.

 The `vdb_init.sh` and `vdb_operations.sh` require the "jq" command line json parser.

Filename: `vdb_init.sh`

The `vdb_init.sh` supports the start, stop, enable, disable, status, and delete command line options.

```
$ ./vdb_init.sh something VBITT
. . .
Unknown option (start | stop | enable | disable | status | delete): something
```

```

Exiting ...
$ ./vdb_init.sh status VBITT
database container reference: ORACLE_DB_CONTAINER-121
source reference: ORACLE_VIRTUAL_SOURCE-25
Runtime Status: "INACTIVE"
Enabled: true
Done ...

```

```

$ ./vdb_init.sh start VBITT
database container reference: ORACLE_DB_CONTAINER-121
source reference: ORACLE_VIRTUAL_SOURCE-25
Job: JOB-894
Current status as of Wed Sep 7 16:04:04 EDT 2016 : RUNNING 0% Completed
Current status as of Wed Sep 7 16:04:14 EDT 2016 : RUNNING 25% Completed
Current status as of Wed Sep 7 16:04:24 EDT 2016 : RUNNING 45% Completed
Job: JOB-894 COMPLETED 100% Completed ...
Done ...

```

```

$ ./vdb_init.sh delete VBITT
database container reference: ORACLE_DB_CONTAINER-123
source reference: ORACLE_VIRTUAL_SOURCE-27
vendor source: OracleVirtualSource
delete parameters type: OracleDeleteParameters
Job: JOB-927
Current status as of Sat Sep 10 12:55:32 EDT 2016 : RUNNING 0% Completed
Current status as of Sat Sep 10 12:55:32 EDT 2016 : RUNNING 0% Completed
Job: JOB-927 COMPLETED 100% Completed ...
Done ...

```

VDB Operations (sync, refresh, rollback)

This script is used to perform a sync (snapshot), refresh, or rollback (reset) on the Delphix Engine source object. All these work on a virtual databases (VDBs), but only a sync operation can be used on dSources.

 The `vdb_init.sh` and `vdb_operations.sh` require the "jq" command line json parser.

Filename: `vdb_operations.sh`

```

$ ./vdb_operations.sh sync VBITT
Session and Login Successful ...
database container reference: ORACLE_DB_CONTAINER-131
provision source container: ORACLE_DB_CONTAINER-129
json> {

```

```

    "type": "OracleSyncParameters"
  }
Job: JOB-998
Current status as of Wed Sep 14 17:05:00 EDT 2016 : RUNNING 0% Completed

```

```

Current status as of Wed Sep 14 17:05:24 EDT 2016 : RUNNING 97% Completed
Job: JOB-998 COMPLETED 100% Completed ...
Done ...

```

Rollback rewinds the virtual database back to the last point in time within the source TimeFlow.

```

$ ./vdb_operations.sh rollback VBITT
Session and Login Successful ...
database container reference: ORACLE_DB_CONTAINER-131
provision source container: ORACLE_DB_CONTAINER-129
json> {
  "type": "OracleRollbackParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "ORACLE_DB_CONTAINER-131"
  }
}
Job: JOB-1000
Current status as of Wed Sep 14 17:06:33 EDT 2016 : RUNNING 0% Completed
Current status as of Wed Sep 14 17:06:43 EDT 2016 : RUNNING 0% Completed
Current status as of Wed Sep 14 17:06:53 EDT 2016 : RUNNING 34% Completed
Current status as of Wed Sep 14 17:07:03 EDT 2016 : RUNNING 34% Completed
Current status as of Wed Sep 14 17:07:13 EDT 2016 : RUNNING 58% Completed
Current status as of Wed Sep 14 17:07:23 EDT 2016 : RUNNING 67% Completed
Current status as of Wed Sep 14 17:07:33 EDT 2016 : RUNNING 70% Completed
Current status as of Wed Sep 14 17:07:43 EDT 2016 : RUNNING 70% Completed
Current status as of Wed Sep 14 17:07:53 EDT 2016 : RUNNING 71% Completed
Current status as of Wed Sep 14 17:08:03 EDT 2016 : RUNNING 72% Completed
Current status as of Wed Sep 14 17:08:13 EDT 2016 : RUNNING 73% Completed
Current status as of Wed Sep 14 17:08:23 EDT 2016 : RUNNING 94% Completed
Current status as of Wed Sep 14 17:08:33 EDT 2016 : RUNNING 96% Completed
Current status as of Wed Sep 14 17:08:43 EDT 2016 : RUNNING 96% Completed
Current status as of Wed Sep 14 17:08:53 EDT 2016 : RUNNING 96% Completed
Job: JOB-1000 COMPLETED 100% Completed ...
Done ...

```

Refresh recreates the virtual database to the last point in time in the respective parent, provision source TimeFlow.

```
$ ./vdb_operations.sh refresh VBITT
Session and Login Successful ...
database container reference: ORACLE_DB_CONTAINER-131
provision source container: ORACLE_DB_CONTAINER-129
json> {
  "type": "OracleRefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "ORACLE_DB_CONTAINER-129"
  }
}
Job: JOB-1005
Current status as of Wed Sep 14 17:10:11 EDT 2016 : RUNNING 0% Completed
Current status as of Wed Sep 14 17:10:21 EDT 2016 : RUNNING 0% Completed
Current status as of Wed Sep 14 17:10:31 EDT 2016 : RUNNING 34% Completed
Current status as of Wed Sep 14 17:10:41 EDT 2016 : RUNNING 35% Completed
Current status as of Wed Sep 14 17:10:51 EDT 2016 : RUNNING 70% Completed
Current status as of Wed Sep 14 17:11:01 EDT 2016 : RUNNING 70% Completed
Current status as of Wed Sep 14 17:11:20 EDT 2016 : RUNNING 72% Completed
Current status as of Wed Sep 14 17:11:31 EDT 2016 : RUNNING 73% Completed
Current status as of Wed Sep 14 17:11:41 EDT 2016 : RUNNING 73% Completed
Job: JOB-1005 COMPLETED 100% Completed ...
Done ...
```

Delphix Self-Service Use Cases for APIs

Create Delphix Self-Service Template

 Jet Stream is now known as Delphix Self-Service.

Filename: *jetstream_template.sh* or *jetstream_template_jq.sh*

Edit the file to update the parameters as required for your environment.

Required for Delphix Self-Service Template ...

```
TPL_NAME="jstpl"           # JetStream Template Name
DATASOURCE_NAME="jsds"    # JetStream Data Source Name
DATASOURCE_VDB="VBITT"   # JetStream Data Source VDB or dSource
```

Sample Output

```
$ ./jetstream_template.sh# or ./Jetstream_template_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
Getting Database Container Reference Value ...
container reference: ORACLE_DB_CONTAINER-45
Create JetStream Template jstpl with Data Source DB VBITT ...
Database: {"type":"OKResult","status":"OK","result":"JS_DATA_TEMPLATE-3","job":null,"
action":"ACTION-547"}

Done ... (no job required for this action)
```

Create Delphix Self-Service Data Container

Filename: *jetstream_container.sh#* or *jetstream_container_jq.sh*

Edit the file to update the parameters as required for your environment.

Required for Delphix Self-Service Container ...

```
TPL_NAME="jstpl"           # JetStream Template Name
DS_NAME="jsds"            # JetStream Data Source Name

DC_NAME="jsdc"            # JetStream Data Container Name
DC_VDB="VBITT2"          # JetStream Data Container VDB
```

Sample Output

```
$ ./jetstream_container.sh# or ./jetstream_container_jq.sh
```

```

Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
Getting Database Container Reference Value ...
container reference: ORACLE_DB_CONTAINER-46
JetStream Data Template: JS_DATA_TEMPLATE-4
JetStream sourceDataLayout: JS_DATA_TEMPLATE-4
Create JetStream Container jsdc with Data Source DB VBITT2 ...
JetStream Data Container Creation Results: {"type":"OKResult","status":"OK","result":"
JS_DATA_CONTAINER-4","job":"JOB-240","action":"ACTION-569"}
Job: JOB-240
Current status as of Wed Aug 17 04:11:54 EDT 2016 : RUNNING 0.0% Completed
Current status as of Wed Aug 17 04:11:54 EDT 2016 : RUNNING 0.0% Completed
Current status as of Wed Aug 17 04:12:04 EDT 2016 : RUNNING 5.0% Completed
Current status as of Wed Aug 17 04:12:14 EDT 2016 : RUNNING 5.0% Completed
Current status as of Wed Aug 17 04:12:24 EDT 2016 : RUNNING 30.0% Completed
Current status as of Wed Aug 17 04:12:34 EDT 2016 : RUNNING 31.0% Completed
Current status as of Wed Aug 17 04:12:44 EDT 2016 : RUNNING 53.0% Completed
Current status as of Wed Aug 17 04:12:54 EDT 2016 : RUNNING 57.0% Completed
Current status as of Wed Aug 17 04:13:04 EDT 2016 : RUNNING 57.0% Completed
Current status as of Wed Aug 17 04:13:14 EDT 2016 : RUNNING 57.0% Completed
Current status as of Wed Aug 17 04:13:24 EDT 2016 : RUNNING 59.0% Completed
Current status as of Wed Aug 17 04:13:34 EDT 2016 : RUNNING 60.0% Completed
Current status as of Wed Aug 17 04:13:44 EDT 2016 : RUNNING 77.0% Completed
Current status as of Wed Aug 17 04:13:54 EDT 2016 : RUNNING 77.0% Completed
Current status as of Wed Aug 17 04:14:04 EDT 2016 : RUNNING 77.0% Completed
Current status as of Wed Aug 17 04:14:14 EDT 2016 : RUNNING 77.0% Completed
Job: JOB-240 COMPLETED 100.0% Completed ...

Done ...

```

Create Delphix Self-Service Bookmark

Filename: *jetstream_api_examples.txt (part 1)*

Create Bookmark ...

Change parameters as required and desired.

```

curl -X POST -k --data @-
http://172.16.160.177/resources/json/delphix/jetstream/bookmark \
-b cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "JSBookmarkCreateParameters",
  "bookmark": {
    "type": "JSBookmark",
    "name": "aalen",
    "branch": "JS_BRANCH-5",
    "shared": false,
    "tags": [
      "A",
      "B",
      "C"
    ]
  }
}

```

```

    ]
  },
  "timelinePointParameters": {
    "type": "JSTimelinePointLatestTimeInput",
    "sourceDataLayout": "JS_DATA_CONTAINER-2"
  }
}
EOF
{"type":"OKResult","status":"OK","result":"JS_BOOKMARK-5","job":"JOB-512","action":
"ACTION-921"}

```

 The timelinePointParameters type "JSTimelinePointLatestTimeInput" is the last point / latest time in the branch!

Filename: *jetstream_bookmark.sh* or *jetstream_bookmark_jq.sh*

Edit the file to update the parameters as required for your environment.

```
DT=`date '+%Y%m%d%H%M%S'`
```

Required for Delphix Self-Service Bookmark ...

```

JS_BRANCH="default"           # JetStream Branch
BM_NAME="aalen_${DT}"        # JetStream Bookmark Name appended timestamp
SHARED="false"               # Share Bookmark true/false
TAGS='"API","Created"'       # Tags Array Values

```

Sample Output

```

$ ./jetstream_bookmark.sh # or ./jetstream_bookmark_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
Getting Jetstream Branch Reference Value ...
branch reference: JS_BRANCH-7
dataLayout container reference: JS_DATA_CONTAINER-4
JetStream Bookmark Creation Results: {"type":"OKResult","status":"OK","result":"JS_BO
OKMARK-4","job":"JOB-251","action":"ACTION-591"}
Job: JOB-251
Current status as of Wed Aug 17 04:59:53 EDT 2016 : COMPLETED 100.0% Completed
Job: JOB-251 COMPLETED 100.0% Completed ...

Done ...

```

Delphix Self-Service Refresh

Filename: *jetstream_api_examples.txt (part 2)*

Use CLI command to get Delphix Self-Service Container Reference

```
/jetstream/container/list

...
"reference": " JS_DATA_CONTAINER-4 ",
"namespace": null,
"name": " jsdc ",
...
```

Refresh Container Information ...

```
=== POST /resources/json/delphix/jetstream/container/ JS_DATA_CONTAINER-4 /refresh

curl -X POST -k --data @http://172.16.160.177/resources/json/delphix/jetstream/
container/ JS_DATA_CONTAINER-4 /refresh \ -b cookies.txt -H "Content-Type:
application/json" <<EOF
{}
EOF

=== RESPONSE ===
{
  "type": "OKResult",
  "status": "OK",
  "result": "",
  "job": "JOB-514",
  "action": "ACTION-924"
}
=== END ===
```

Filename: *jetstream_refresh.sh* or *jetstream_refresh_jq.sh*

Edit the file to update the parameters as required for your environment.

Required for Delphix Self-Service Refresh ...

```
CONTAINER_NAME="jsdc"          # Jetstream Container Name
```

Sample Output

```
$ ./jetstream_refresh.sh # or ./jetstream_refresh_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
```

```
Session and Login Successful ...
Getting Jetstream Container Reference Value ...
container reference: JS_DATA_CONTAINER-4
abitterman-mbpro:JetStream abitterman$ vi jetstream_refresh.sh
abitterman-mbpro:JetStream abitterman$ ./jetstream_refresh.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
Getting Jetstream Container Reference Value ...
container reference: JS_DATA_CONTAINER-4
JetStream Refresh API Results: {"type":"OKResult","status":"OK","result":"","job":"JOB-257","action":"ACTION-602"}
Job: JOB-257
Current status as of Wed Aug 17 05:13:15 EDT 2016 : RUNNING 2.0% Completed
Current status as of Wed Aug 17 05:13:15 EDT 2016 : RUNNING 2.0% Completed
Current status as of Wed Aug 17 05:13:25 EDT 2016 : RUNNING 5.0% Completed
Current status as of Wed Aug 17 05:13:35 EDT 2016 : RUNNING 5.0% Completed
Current status as of Wed Aug 17 05:13:45 EDT 2016 : RUNNING 30.0% Completed
Current status as of Wed Aug 17 05:13:55 EDT 2016 : RUNNING 42.0% Completed
Current status as of Wed Aug 17 05:14:05 EDT 2016 : RUNNING 55.0% Completed
Current status as of Wed Aug 17 05:14:15 EDT 2016 : RUNNING 58.0% Completed
Current status as of Wed Aug 17 05:14:25 EDT 2016 : RUNNING 58.0% Completed
Current status as of Wed Aug 17 05:14:35 EDT 2016 : RUNNING 58.0% Completed
Current status as of Wed Aug 17 05:14:45 EDT 2016 : RUNNING 60.0% Completed
Current status as of Wed Aug 17 05:14:55 EDT 2016 : RUNNING 62.0% Completed
Current status as of Wed Aug 17 05:15:05 EDT 2016 : RUNNING 77.0% Completed
Current status as of Wed Aug 17 05:15:15 EDT 2016 : RUNNING 77.0% Completed
Job: JOB-257 COMPLETED 100.0% Completed ...

Done ...
```

Masking Use Cases

Masking API Client

The Delphix Masking Engine now features an interactive API client that can generate commands specific to your masking engine. With those commands, you can:

- make changes to your engine
- copy and paste the commands to write code that can automate your masking activities



The API client will make real changes to your virtual machine. Any operations you run using the API Client will persist on the machine!

To access the Masking API client, use the following URL: <http://myMaskingEngine.com/masking/api-client/>, replacing "myMaskingEngine.com" with the hostname or IP address of your virtual machine.

For detailed examples of using API calls to automate masking, see the [Masking APIs](#).

Masking in Parallel

Delphix masking supports launching masking jobs in parallel. When jobs have no dependencies, you can initiate parallel masking API jobs (with wrapper code as required) to allow the jobs to be run as a pre and/or post hook.

Oracle Use Cases for APIs

Oracle Link + Snapshot (Sync)

The following script ingests links an environment database dSource (Oracle SID / Instance) and then takes a snapshot. See parameters for required values that you must provide.

This script demonstrates how to use name values inputs and get the respective Delphix object and/or object reference for use in the json input in downstream API calls.

Filename: link_oracle.sh # or link_oracle_jq.sh

Edit the file to update the parameters as required for your environment.

```
#####
#Parameter Initialization

DMIP=172.16.160.195
#DMPORT=8282
DMUSER=delphix_admin
DMPASS=delphix
. . .

Required for Database Link and Sync ...

SOURCE_SID="DPXDEV01"           # Source Environment Database SID
SOURCE_NAME="DPXDEV01"         # Delphix dSource Name
SOURCE_ENV="Oracle Target"     # Source Environment Name
SOURCE_GRP="Oracle_Source"     # Delphix Group Name
DB_USER="delphixdb"           # Source Database SID user account
DB_PASS="delphixdb"           # Source Database SID user password

#####
# NO CHANGES REQUIRED BELOW THIS POINT #
#####

$ ./link_oracle.sh      # or   ./link_oracle_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
group reference: GROUP-35
sourceconfig reference: ORACLE_SINGLE_CONFIG-1
primaryUser reference: HOST_USER-3
Linking Source Database ...
Job: JOB-92
Job: JOB-92 100.0% Completed ...
Container: ORACLE_DB_CONTAINER-19
Running SnapSync ...
Job: JOB-93
Current status as of Mon Aug 15 13:07:53 EDT 2016 : RUNNING : 0.0% Completed
Current status as of Mon Aug 15 13:08:03 EDT 2016 : RUNNING : 15.0% Completed
Current status as of Mon Aug 15 13:08:13 EDT 2016 : RUNNING : 35.0% Completed
```

```
Current status as of Mon Aug 15 13:08:24 EDT 2016 : RUNNING : 59.0% Completed
```

```
Current status as of Mon Aug 15 13:08:34 EDT 2016 : RUNNING : 66.0% Completed
Current status as of Mon Aug 15 13:08:44 EDT 2016 : RUNNING : 74.0% Completed
Job: JOB-93 100.0% Completed ...
Done ...
$
```

Oracle Provision

Filename: provision_oracle.txt

Shown below is how to use the CLI to provision an Oracle 11g database that is already ingested into the Delphix Engine.

The key is to get the object reference names first. For example, to get the source database container name:

```
ssh delphix_admin[delphix_engine_ip_address_or_hostname]
> database
> ls
...
> select "[database_name]"
> ls
Delphix5002HWv7 database> select 'DPXDEV01'

Delphix5002HWv7 database 'DPXDEV01'> ls
Properties
  type: OracleDatabaseContainer
  name: DPXDEV01
  . . .
  reference: ORACLE_DB_CONTAINER-18
  . . .
```

Minimum parameters required to provision:

```
Delphix5002HWv7 database provision > *commit
=== POST /resources/json/delphix/database/provision ===
{
  "type": "OracleProvisionParameters",
  "container": {
    "type": "OracleDatabaseContainer",
    "name": "VBITT" , # Delphix Object Name, Typically matches VDB name
    "group": "GROUP-36" # group ls select "[group_name]" ls
  },
  "source": {
    "type": "OracleVirtualSource",
    "mountBase": "/mnt/provision" # Delphix Filesystem Mount path
  },
  "sourceConfig": {
```

```

    "type": "OracleSIConfig",
    "repository": "ORACLE_INSTALL-3" , # repository, select "[repository_name]"
    "databaseName": "VBITT" , # New VDB Name
    "uniqueName": "VBITT",
    "instance": {
        "type": "OracleInstance",
        "instanceName": "VBITT",
        "instanceNumber": 1
    }
},
"timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "ORACLE_DB_CONTAINER-18" # select "[database_name]" ls
}
}
=== RESPONSE ===

```

Sample CLI Session

```

Delphix5002HWv7 database> setopt trace=false
Delphix5002HWv7 database> provision
Delphix5002HWv7 database provision > *ls
Properties
  type: OracleProvisionParameters
  container:
    type: OracleDatabaseContainer
    name: (required)
    description: (unset)
    diagnoseNoLoggingFaults: true
    group: (required)
    performanceMode: DISABLED
    preProvisioningEnabled: false
    sourcingPolicy: (unset)
  credential: (unset)
  maskingJob: (unset)
  newDBID: false
  openResetlogs: true
  physicalStandby: false
  source:
    type: OracleLiveSource
    name: (unset)
    archivelogMode: true
    config: (unset)
    configParams: (unset)
    configTemplate: (unset)
    customEnvVars: (unset)
    dataAgeWarningThreshold: 900sec
    fileMappingRules: (unset)
    manualProvisioning: false
    mountBase: (required)
    nodeListenerList: (unset)
    operations: (unset)

```

```

    redoLogGroups: 3
    redoLogSizeInMB: 0
sourceConfig:
  type: OraclePDBConfig
  cdbConfig: (required)
  databaseName: (required)
  environmentUser: (unset)
  linkingEnabled: true
  repository: (unset)
  services: (unset)
timeflowPointParameters:
  type: TimeflowPointSemantic
  container: (required)
  location: LATEST_POINT
username: (unset)
Operations
defaults
Delphix5002HWv7 database provision > *edit container
Delphix5002HWv7 database provision container> *ls
Properties
  type: OracleDatabaseContainer
  name: (required)
  description: (unset)
  diagnoseNoLoggingFaults: true
  group: (required)
  performanceMode: DISABLED
  preProvisioningEnabled: false
  sourcingPolicy: (unset)
Delphix5002HWv7 database provision container> *set name=VBITT
Delphix5002HWv7 database provision container> *set group=GROUP-36
Delphix5002HWv7 database provision container> *back
Delphix5002HWv7 database provision > *edit source
Delphix5002HWv7 database provision source > *ls
Properties
  type: OracleLiveSource
  name: (unset)
  archiveLogMode: true
  config: (unset)
  configParams: (unset)
  configTemplate: (unset)
  customEnvVars: (unset)
  dataAgeWarningThreshold: 900sec
  fileMappingRules: (unset)
  manualProvisioning: false
  mountBase: (required)
  nodeListenerList: (unset)
  operations: (unset)
  redoLogGroups: 3
  redoLogSizeInMB: 0
Delphix5002HWv7 database provision source > *set type=OracleVirtualSource
Delphix5002HWv7 database provision source > *set mountBase=/mnt/provision
Delphix5002HWv7 database provision source > *back
Delphix5002HWv7 database provision > *edit sourceConfig

```

```

Delphix5002HWv7 database provision sourceConfig > *ls
Properties
  type: OraclePDBConfig
  cdbConfig: (required)
  databaseName: (required)
  environmentUser: (unset)
  linkingEnabled: true
  repository: (unset)
  services: (unset)
Delphix5002HWv7 database provision sourceConfig > *set type=OracleSIConfig
Delphix5002HWv7 database provision sourceConfig > *ls
Properties
  type: OracleSIConfig
  databaseName: (required)
  environmentUser: (unset)
  instance: (required)
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: (required)
  services: (unset)
  uniqueName: (required)
Delphix5002HWv7 database provision sourceConfig > *set databaseName=VBITT
Delphix5002HWv7 database provision sourceConfig > *set repository=ORACLE_INSTALL-3
Delphix5002HWv7 database provision sourceConfig > *set uniqueName=VBITT
Delphix5002HWv7 database provision sourceConfig > *set instance.instanceName=VBITT
Delphix5002HWv7 database provision sourceConfig > *set instance.instanceNumber=1
Delphix5002HWv7 database provision sourceConfig > *ls
Properties
  type: OracleSIConfig
  databaseName: VBITT
  environmentUser: (unset)
  instance:
    type: OracleInstance
    instanceName: VBITT
    instanceNumber: 1
  linkingEnabled: true
  nonSysCredentials: (unset)
  nonSysUser: (unset)
  repository: '/u02/ora/app/product/11.2.0/dbhome_1'
  services: (unset)
  uniqueName: VBITT
Delphix5002HWv7 database provision sourceConfig > *back
Delphix5002HWv7 database provision > *edit timeflowPointParameters
Delphix5002HWv7 database provision timeflowPointParameters> *ls
Properties
  type: TimeflowPointSemantic
  container: (required)
  location: LATEST_POINT
Delphix5002HWv7 database provision timeflowPointParameters> *set
container=ORACLE_DB_CONTAINER-18
Delphix5002HWv7 database provision timeflowPointParameters> *back
Delphix5002HWv7 database provision > *commit

```

```

VBITT
Dispatched job JOB-348
DB_PROVISION job started for "Oracle Target Virtual Databases/VBITT".
Starting provision of the virtual database "VBITT".
Creating new TimeFlow.
Generating recovery scripts.
Exporting storage.
Mounting filesystems for the virtual database instance "1".
Mounting read-only archive log filesystem for the virtual database instance "1".
Recovering Oracle database.
\|/-
Opening the virtual database "VBITT".
Opening Oracle database.
Oracle recovery was successful.
Unmounting read-only archive log filesystem for the virtual database instance "1".
The virtual database "VBITT" was successfully provisioned.
DB_PROVISION job for "Oracle Target Virtual Databases/VBITT" completed
successfully.
Delphix5002HWv7 database>

```

With the **setopt trace=true** option set, you can convert the JSON output from the above CLI provision command to the RESTful API cURL commands. If VBITT exists, be sure to delete it first.

Request:

```

curl X POST -k --data @http://172.16.160.177/resources/json/delphix/database/
provision \
  -b cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "OracleProvisionParameters",
  "container": {
    "type": "OracleDatabaseContainer",
    "name": "VBITT",
    "group": "GROUP-36"
  },
  "source": {
    "type": "OracleVirtualSource",
    "mountBase": "/mnt/provision"
  },
  "sourceConfig": {
    "type": "OracleSIConfig",
    "repository": "ORACLE_INSTALL-3",
    "databaseName": "VBITT",
    "uniqueName": "VBITT",
    "instance": {
      "type": "OracleInstance",
      "instanceName": "VBITT",
      "instanceNumber": 1
    }
  },
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",

```

```

        "container": "ORACLE_DB_CONTAINER-18"
    }
}
EOF

```

Response:

```

{"type":"OKResult","status":"OK","result":"ORACLE_DB_CONTAINER-22","job":"JOB-353",
"action":"ACTION-649"}

```

Put all the commands above within a shell script to automate the complete process of provisioning an Oracle 11.2.0.4 database.

⚠ Notice that the script below looks up 4 object references for use within the JSON input into the API.

Filename: provision_oracle.sh# or provision_oracle_jq.sh

Edit the file to update the parameters as required for your environment.

```

#####
#   DELPHIX CORP   #
#####
#Parameter Initialization

DMIP=172.16.160.195
DMUSER=delphix_admin
DMPASS=delphix
COOKIE=~/.cookies.txt
COOKIE=`eval echo $COOKIE`
CONTENT_TYPE="Content-Type: application/json"
DELAYTIMESEC=10
BaseUrl="http://${DMIP}/resources/json/delphix"
#
Required for Database Link and Sync ...
#
VDB_NAME="VBITT"           # Delphix VDB Name
MOUNT_BASE="/mnt/provision" # Delphix Engine Mount Path
SOURCE_GRP="Oracle_Target" # Delphix Engine Group Name
TARGET_ENV="Oracle Target" # Target Environment used to get repository reference
value
SOURCE_SID="DPXDEV01"     # dSource name used to get db container reference
value
#####
#   NO CHANGES REQUIRED BELOW THIS POINT   #
#####

```

Sample Output

```

$ ./provision_oracle.sh# or ./provision_oracle_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
group reference: GROUP-36
container reference: ORACLE_DB_CONTAINER-36
env reference: UNIX_HOST_ENVIRONMENT-3
repository reference: ORACLE_INSTALL-1
Provisioning VDB from Source Database ...
Job: JOB-155
Current status as of Mon Aug 15 23:40:51 EDT 2016 : RUNNING 0.0% Completed
Current status as of Mon Aug 15 23:40:51 EDT 2016 : RUNNING 0.0% Completed
Current status as of Mon Aug 15 23:41:01 EDT 2016 : RUNNING 9.0% Completed
Current status as of Mon Aug 15 23:41:11 EDT 2016 : RUNNING 45.0% Completed
Current status as of Mon Aug 15 23:41:21 EDT 2016 : RUNNING 45.0% Completed
Current status as of Mon Aug 15 23:41:31 EDT 2016 : RUNNING 46.0% Completed
Current status as of Mon Aug 15 23:41:41 EDT 2016 : RUNNING 48.0% Completed
Current status as of Mon Aug 15 23:41:51 EDT 2016 : RUNNING 60.0% Completed
Job: JOB-155 COMPLETED 100.0% Completed ...
Done ...
$

```

Filename: provision_oracle_child.sh# or provision_oracle_child_jq.sh

```

$ ./provision_oracle_child.sh# or ./provision_oracle_child_jq.sh
Authenticating on http://172.16.160.195/resources/json/delphix
Session and Login Successful ...
group reference: GROUP-36
container reference: ORACLE_DB_CONTAINER-118
env reference: UNIX_HOST_ENVIRONMENT-9
repository reference: ORACLE_INSTALL-6
Provisioning VDB from Source Database ...
Job: JOB-857
Current status as of Mon Sep 5 22:48:28 EDT 2016 : RUNNING 0.0% Completed
Current status as of Mon Sep 5 22:48:28 EDT 2016 : RUNNING 0.0% Completed
Current status as of Mon Sep 5 22:48:38 EDT 2016 : RUNNING 9.0% Completed
Current status as of Mon Sep 5 22:48:48 EDT 2016 : RUNNING 27.0% Completed
Current status as of Mon Sep 5 22:48:58 EDT 2016 : RUNNING 42.0% Completed
Current status as of Mon Sep 5 22:49:08 EDT 2016 : RUNNING 45.0% Completed
Current status as of Mon Sep 5 22:49:28 EDT 2016 : RUNNING 46.0% Completed
Current status as of Mon Sep 5 22:49:38 EDT 2016 : RUNNING 48.0% Completed
Current status as of Mon Sep 5 22:49:48 EDT 2016 : RUNNING 51.0% Completed
Current status as of Mon Sep 5 22:50:08 EDT 2016 : RUNNING 71.0% Completed
Job: JOB-857 COMPLETED 100.0% Completed ...
Done ...

```

SQL Server API Use Cases

SQL Server Link/Ingest Environment dSource

For the **Window Target** environment, the dSource **delphixdb** in **MSSQLSERVER** instance will be linked/ingested into the Delphix Engine. It will appear in the **Windows_Source** group below.

Filename: *link_sqlserver.ps1*

```
PS> . .\link_sqlserver.ps1
Authenticating on http://172.16.160.195/resources/json/delphix
Login Successful ...
group reference: GROUP-34
sourceconfig reference: MSSQL_SINGLE_CONFIG-26
env reference: WINDOWS_HOST_ENVIRONMENT-7
repository reference: MSSQL_INSTANCE-4
database link API Results: {"type":"OKResult","status":"OK","result":"MSSQL_DB_CONTAI
NER-114","job":"JOB-819","action":
"ACTION-1659"}
DB Container: MSSQL_DB_CONTAINER-114
```

```
Job # JOB-819
***** waiting for status *****
Current status as of 09/05/2016 11:41:13 : COMPLETED : 100.0% Completed
Job COMPLETED Successfully.
```

```
JOB JOB-820
waiting for status *****
Current status as of 09/05/2016 11:41:23 : RUNNING : 5.0% Completed
Current status as of 09/05/2016 11:41:44 : RUNNING : 9.0% Completed
```

```
Current status as of 09/05/2016 11:41:54 : RUNNING : 56.0% Completed
Job COMPLETED Successfully.
Done ...
```

Successful dSource linked/ingested into the Delphix Engine.

SQL Server Provision

The example below is done from the command line once you know the parameters and reference object names.

Filename: *windows_sqlserver_provision.txt*

Create these 3 JSON text files:

```

session.json
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 5,
    "micro": 3
  }
}
login.json
{
  "type": "LoginRequest",
  "username": "delphix_admin",
  "password": "delphix"
}
provision.json
{
  "type": "MSSqlProvisionParameters",
  "container": {
    "type": "MSSqlDatabaseContainer",
    "name": "Vbitt00",
    "group": "GROUP-36",
    "sourcingPolicy": {
      "type": "SourcingPolicy",
      "loadFromBackup": false,
      "logsyncEnabled": false
    },
    "validatedSyncMode": "TRANSACTION_LOG"
  },
  "source": {
    "type": "MSSqlVirtualSource",
    "operations": {
      "type": "VirtualSourceOperations",
      "configureClone": [],
      "postRefresh": [],
      "postRollback": [],
      "postSnapshot": [],
      "preRefresh": [],
      "preSnapshot": []
    }
  },
  "sourceConfig": {
    "type": "MSSqlSIConfig",
    "linkingEnabled": false,
    "repository": "MSSQL_INSTANCE-1",
    "databaseName": "Vbitt00",
    "recoveryModel": "SIMPLE",
    "instance": {
      "type": "MSSqlInstanceConfig",
      "host": "WINDOWS_HOST-1"
    }
  }
}

```

```

    },
    "timeflowPointParameters": {
        "type": "TimeflowPointSemantic",
        "container": "MSSQL_DB_CONTAINER-23",
        "location": "LATEST_SNAPSHOT"
    }
}

```

This works on Windows Powershell Command Prompt

 Use curl, curl.exe or modify the default alias.

```

curl --insecure -c cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@session.json" http://172.16.160.153/resources/json/delphix/session

curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@login.json" http://172.16.160.153/resources/json/delphix/login

curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@provision.json" http://172.16.160.153/resources/json/delphix/database/provision

```

Plug in the returned JOB #

```

curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k http://
/172.16.160.153/resources/json/delphix/notification?channel=JOB-428

```

Get Example

```

curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k http://
/172.16.160.153/resources/json/delphix/system

```

Complete example.

Provision the newly created **delphixdb** dSource in the **Windows_Source** group to a virtual database VBITT in the **Windows_Target** group.

Filename: provision_sqlserver.ps1

Variables...

```

$nl = [Environment]::NewLine
$BaseUrl = " http://172.16.160.195/resources/json/delphix "
$cookie = "cookies.txt"

```

```
$delphix_user = "delphix_admin"
$delphix_pass = "delphix"
. . .
```

Required for Provisioning Virtual Database ...

```
$$SOURCE_SID="delphixdb"          # dSource name used to get db container reference
value

$VDB_NAME="VBITT"                 # Delphix VDB Name
$TARGET_GRP="Windows_Target"     # Delphix Engine Group Name
$TARGET_ENV="Window Target"      # Target Environment used to get repository reference
value
$TARGET_REP="MSSQLSERVER"        # Target Environment Repository / Instance name

#####
# NO CHANGES REQUIRED BELOW THIS POINT #
#####
```

Sample Run Output

```
PS> . .\provision_sqlserver.ps1
Authenticating on http://172.16.160.195/resources/json/delphix
Login Successful ...
group reference: GROUP-37
container reference: MSSQL_DB_CONTAINER-114
env reference: WINDOWS_HOST_ENVIRONMENT-7
repository reference: MSSQL_INSTANCE-4
database provision API Results: {"type":"OKResult","status":"OK","result":"MSSQL_DB_C
ONTAINER-115","job":"JOB-822","act
ion":"ACTION-1664"}
DB Container: MSSQL_DB_CONTAINER-115

Job # JOB-822

jobState RUNNING
percentComplete 0.0
***** waiting for status *****
Current status as of 09/05/2016 11:43:51 : RUNNING : 0.0% Completed
Current status as of 09/05/2016 11:44:01 : RUNNING : 3.0% Completed
Current status as of 09/05/2016 11:44:12 : RUNNING : 11.0% Completed
Current status as of 09/05/2016 11:44:22 : RUNNING : 18.0% Completed
Current status as of 09/05/2016 11:44:32 : RUNNING : 18.0% Completed
```

```
Current status as of 09/05/2016 11:44:52 : RUNNING : 75.0% Completed
```

```
Job COMPLETED Successfully.
```

```
Done ...
```

SQL Server Refresh

The following are curl commands that can be issued from the Powershell command line. For inclusion within a Powershell script, see the masking example, *masking.ps1*.

Filename: *windows_sqlserver_refresh.txt*

MS SQL Server Refresh Example ...

Session ...

```
curl --insecure -c cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@session.json" http://172.16.160.179/resources/json/delphix/session
```

Filename: session.json

```
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 5,
    "micro": 3
  }
}
```

```
PS> *curl --insecure -c cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@session.json" http://172.16.160.179/resources/json/delphix/session*
HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Set-Cookie: JSESSIONID=8DE0362F5BBD73E6BFA9E13FF111E78C; Path=/resources/; HttpOnly
Content-Type: application/json
Content-Length: 179
Date: Thu, 16 Jun 2016 07:24:34 GMT
```

```
{"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"API
Version","major":1,"minor":5,"micro":3},"locale":null,"client":null},"job":null,
"action":null}
```

```
PS>
```

Login ...

```
curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@login.json" http://172.16.160.179/resources/json/delphix/login
```

Filename:login.json

```
{
  "type": "LoginRequest",
  "username": "delphix_admin",
  "password": "delphix"
}
```

```
PS> *curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@login.json" http://172.16.160.179/resources/json/delphix/login*
HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Content-Type: application/json
Content-Length: 76
Date: Thu, 16 Jun 2016 07:25:39 GMT
```

```
{"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
PS C:\Users\Administrator>
```

List Databases ...

```
curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k http://
/172.16.160.179/resources/json/delphix/database
```

```
PS> *curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k
http://172.16.160.179/resources/json/delphix/database*
HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Content-Type: application/json
Content-Length: 4062
Date: Thu, 16 Jun 2016 07:27:14 GMT
```

```
{"type":"ListResult","status":"OK","result":[...
...
{"type":"MSSqlDatabaseContainer","reference":"MSSQL_DB_CONTAINER-37","namespace":null
,"name":"Vdelphix_demo","group":"GR
OUP-35","provisionContainer":"MSSQL_DB_CONTAINER-36","creationTime":"2016-06-16T07:09
:06.222Z","currentTimeflow":"MSSQL_
TIMEFLOW-38","previousTimeflow":"MSSQL_TIMEFLOW-37","description":null,"runtime":
...
{"type":"So{"type":"MSSqlDatabaseContainer","reference":"MSSQL_DB_CONTAINER-36","
namespace":null,"name":"delphix_demo","
```

```
group":"GROUP-35","provisionContainer":null,"creationTime":"2016-06-16T07:07:49.939Z",
,"currentTimeflow":"MSSQL_TIMEFLOW-36","previousTimeflow":null,"description":"","runtime":
...
...}], "job":null, "action":null, "total":6, "overflow":false}
```

PS>

Need Reference Object from Database Information ...

For Parent Source Database delphix_demo, reference object is MSSQL_DB_CONTAINER-36

For Virtual Database Vdelphix_demo, reference object is MSSQL_DB_CONTAINER-37

[Optional: Get Database Info ...]

```
curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k http://172.16.160.179/resources/json/delphix/database/MSSQL_DB_CONTAINER-37
```

```
PS> *curl --insecure -b cookies.txt -i -X GET -H "Content-Type: application/json" -k
http://172.16.160.179/resources/json/delphix/database/MSSQL_DB_CONTAINER-37*
HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Content-Type: application/json
Content-Length: 696
Date: Thu, 16 Jun 2016 07:35:42 GMT
```

```
{"type":"OKResult","status":"OK","result":{"type":"MSSqlDatabaseContainer","reference":"MSSQL_DB_CONTAINER-37","namespace":null,"name":"Vdelphix_demo","group":"GROUP-35","provisionContainer":"MSSQL_DB_CONTAINER-36","creationTime":"2016-06-16T07:09:06.222Z","currentTimeflow":"MSSQL_TIMEFLOW-38","previousTimeflow":"MSSQL_TIMEFLOW-37","description":null,"runtime":{"type":"MSSqlDBContainerRuntime","logSyncActive":false,"preProvisioningStatus":null,"lastRestoredBackupSetUUID":null},"os":"Windows","processor":"x86","sourcingPolicy":{"type":"SourcingPolicy","logsyncEnabled":false,"loadFromBackup":false},"performanceMode":"DISABLED","delphixManaged":true,"masked":false},"job":null,"action":null}
```

PS>

Refresh Vdelphix_demo using parent delphix_demo (MSSQL_DB_CONTAINER-36) with the latest timecard ...

```
curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d "@refresh.json" http://172.16.160.179/resources/json/delphix/database/MSSQL_DB_CONTAINER-37/refresh
```

```
=== POST /resources/json/delphix/database/MSSQL_DB_CONTAINER-37/refresh ===
refresh.json
{
  "type": "RefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "MSSQL_DB_CONTAINER-36"
  }
}
PS> *curl --insecure -b cookies.txt -i -X POST -H "Content-Type: application/json" -d
"@refresh.json" http://172.16.160.179/resources/json/delphix/database/
MSSQL_DB_CONTAINER-37/refresh*
HTTP/1.1 200 OK
Server: Apache-Coyote/1.1
Content-Type: application/json
Content-Length: 82
Date: Thu, 16 Jun 2016 07:40:44 GMT

{"type":"OKResult","status":"OK","result":"","job":"JOB-60","action":"ACTION-167"}

PS>
```

[Observer Delphix GUI Action]

Done with SQL Server VDB Refresh ...

API Programming Language Examples

The following programming language examples are just to show the bare minimum authentication and a sample functional API call. There are numerous modules, libraries, methods, functions, and code examples to further demonstrate how the languages work with the Delphix APIs and JSON data strings/objects.

You can execute PHP, Perl, and Python languages from the command line and/or from within a Web Server such as Apache or IIS. The following examples are formatted for command line / terminal output.

PHP

PHP provides cURL and JSON modules.

```
$ php -i | grep -iE "cURL|json"
curl
curl support => enabled
curl Information => 7.43.0
json
json support => enabled
json version => 1.2.1
```

Filename: *delphix_curl.php*

Sample Output:

```
$ php -f delphix_curl.php
Session json> {"type":"APISession","version":{"type":"APIVersion","major":1,"minor":7,"micro":0}}
Session Results> {"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"APIVersion","major":1,"minor":7,"micro":0},"locale":null,"client":null},"job":null,"action":null}
Login json> {"type":"LoginRequest","username":"delphix_admin","password":"delphix"}
Login Results> {"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
Calling About API ...
About Results> {"type":"OKResult","status":"OK","result":{"type":"PublicSystemInfo","productType":"standard","productName":"Delphix Engine","buildTitle":"Delphix Engine 5.1.1.0","buildTimestamp":"2016-07-21T07:23:41.000Z","buildVersion":{"type":"VersionInfo","major":5,"minor":1,"micro":1,"patch":0},"configured":true,"enabledFeatures":["XPP","MSSQLH00KS"],"apiVersion":{"type":"APIVersion","major":1,"minor":8,"micro":0},"banner":null,"locales":["en-US"],"currentLocale":"en-US"},"job":null,"action":null}
Converting json string to a PHP Array
stdClass Object
(
    [type] => OKResult
    [status] => OK
    [result] => stdClass Object
        (
            [type] => PublicSystemInfo
            [productType] => standard
```

```

[productName] => Delphix Engine
[buildTitle] => Delphix Engine 5.1.1.0
[buildTimestamp] => 2016-07-21T07:23:41.000Z
[buildVersion] => stdClass Object
(
  [type] => VersionInfo
  [major] => 5
  [minor] => 1
  [micro] => 1
  [patch] => 0
)
[configured] => 1
[enabledFeatures] => Array
(
  [0] => XPP
  [1] => MSSQLHOOKS
)
[apiVersion] => stdClass Object
(
  [type] => APIVersion
  [major] => 1
  [minor] => 8
  [micro] => 0
)
[banner] =>
[locales] => Array
(
  [0] => en-US
)
[currentLocale] => en-US
)
[job] =>
[action] =>
)

```

Perl

Perl provides a couple of methods for working with cURL: operating system calls, WWW::Curl (libcurl) module, or LWP::Curl module. The sample below simply logs into the Delphix Engine and lists the current Delphix Environments.

Filename: [perl_curl.pl](#)

Sample Output:

```

$ perl perl_curl.pl
Testing cURL on Perl ...

Session Results: {"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"APIVersion","major":1,"minor":7,"micro":0},"locale":null,"client":null},"job":null,"action":null}

```

```
Login Results: {"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
```

```
Environment Results: {"type":"ListResult","status":"OK","result":[{"type":"WindowsHostEnvironment","reference":"WINDOWS_HOST_ENVIRONMENT-7","namespace":null,"name":"Window Target","description":null,"primaryUser":"HOST_USER-7","enabled":false,"host":"WINDOWS_HOST-6","proxy":null},{"type":"UnixHostEnvironment","reference":"UNIX_HOST_ENVIRONMENT-9","namespace":null,"name":"Oracle Target","description":"","primaryUser":"HOST_USER-9","enabled":true,"host":"UNIX_HOST-8","aseHostEnvironmentParameters":null}], "job":null,"action":null,"total":2,"overflow":false}
```

Done

Python

Delphix has an extensive resource library for using Python with the Delphix Engine.

[CLI to Python Transition](#)

Delphix python module

Blogs

<https://github.com/CloudSurgeon/delphixpy-examples>

Related Videos

- <https://vimeo.com/164779308>
- <https://vimeo.com/170187276>
- <https://vimeo.com/170896907>

Simple Python program to authenticate and get the "about" API results. This script requires the "request" and "json" modules.

<http://stackoverflow.com/questions/17309288/importerror-no-module-named-requests>

Filename: *auth.py*

Sample Output

```
$ python auth.py
Authenticating URL http://172.16.160.195/resources/json/delphix ...
{"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"API Version","major":1,"minor":7,"micro":0},"locale":null,"client":null},"job":null,"action":null}
Login ...
{"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
About ...
{"type":"OKResult","status":"OK","result":{"type":"PublicSystemInfo","productType":"standard","productName":"Delphix Engine","buildTitle":"Delphix Engine 5.1.1.0","buildTimestamp":"2016-07-21T07:23:41.000Z","buildVersion":{"type":"VersionInfo","major":5,"minor":1,"micro":1,"patch":0},"configured":true,"enabledFeatures":["XPP","MSSQLHOOKS"]}
```

```
, "apiVersion": {"type": "APIVersion", "major": 1, "minor": 8, "micro": 0}, "banner": null, "locales": ["en-US"], "currentLocale": "en-US"}, "job": null, "action": null}
JSON Parsing Examples ...
OK
Delphix Engine 5.1.1.0
1
```

JSP (Java Server Pages)

Java Server Pages are typically used for the web formatting and output, but you can also use JSP for application logic processing and native Java code integration, although this is scorned by the purest and most logical thinking programmers.

Filename: *delphix_http.jsp*

Sample Output:


Browser URL: http://localhost:8080/delphix_http.jsp



```
Trying ...
session> {"type":"OKResult","status":"OK","result":{"type":"APISession","version":{"type":"APIVersion","major":1,"minor":7,"micro":0},"locale":null,"client":null,"job":null,"action":null}}
cookie> JSESSIONID=D16845959B873C64F7460E99A67BFFEB
login> {"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
system results> {"type":"OKResult",
"status":"OK",
"result":{"type":"SystemInfo",
"productType":"standard",
"productName":"Delphix Engine",
"buildTitle":"Delphix Engine 5.1.1.0",
"buildTimestamp":"2016-07-21T07:23:41.000Z",
"buildVersion":{"type":"VersionInfo",
"major":5,
"minor":1,
"micro":1,
"patch":0},
"configured":true,
"enabledFeatures":["XPP",
"MSSQLHOOKS"],
"apiVersion":{"type":"APIVersion",
"major":1,
"minor":8,
"micro":0},
"banner":null,
"locales":["en-US"],
"currentLocale":"en-US",
"hostname":"Delphix5110HWv8",
"sshPublicKey":"ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQDQsp7AjbhFQh9yBq7273B+qtPKmCu1B18nPvr08yjt/IzeM4qKk7caxExQS9rpfU8AWoT7e8ESV7NkBMUzOhrHnLsuJtxPqeYqoeMubVxYj
root@Delphix5110HWv8n",
"memorySize":8.58107904E9,
"platform":"VMware with BIOS date 05/20/2014",
"uuid":"564d7e1d-f4cb-f910-98fd-348d74817683",
"processors":[{"type":"CPUInfo",
"speed":2.5E9,
"cores":1}],
"storageUsed":3.121203712E9,
"storageTotal":2.0673724416E10,
"installationTime":"2016-07-27T13:28:46.000Z"},
"job":null,
"action":null}
```

Java

Java methods and classes allow coding logic to be effectively re-used, extended and modularized for flexible applications. Using the Java code embedded within the JSP file, code is logically placed into respective classes and methods.



API Timeflows

From earlier, the RESTful URL for a virtual database refresh will look like:

http://<delphix_engine>/resources/json/delphix/database/MSSQL_DB_CONTAINER-39/refresh

where the **MSSQL_DB_CONTAINER-39** represents the target virtualized database to refresh and we need to POST the JSON data to the URL for processing.

```
{
  "type": "RefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "MSSQL_DB_CONTAINER-38"
  }
}
```

Timeflow Parameters

The **"timeflowPointParameters"** key has 6 **"type": "..."** options which each have their own set of parameters. The type **"TimeflowPointSemantic"** uses the default LATEST_POINT within the source container. Now, for more on timeflowPointParameters.

http://<delphix_engine>/api/#TimeflowPointParameters

TimeflowPointParameters

Parameters indicating a TimeFlow point to use as input to database operations.

TypedObject

TimeflowPointParameters

Direct Known Subclasses:

TimeflowPointTimestamp, TimeflowPointSnapshot, TimeflowPointSemantic, TimeflowPointLocation, TimeflowPointBookmark, TimeflowPointBookmarkTag

TimeflowPointTimestamp

timeflow	Reference to TimeFlow containing this point. Type: Reference to Timeflow Constraints: Required: true
timestamp	The logical time corresponding to the TimeFlow location. Type: date Constraints: Required: true

TimeflowPointSnapshot

snapshot	Reference to the snapshot. Type: Reference to TimeflowSnapshot
----------	---

TimeflowPointSemantic

Semantic reference to a Timeflow point.

⚠ The reference is relative to a container and not a TimeFlow. If the container contains multiple TimeFlows, the Delphix Engine will evaluate the semantic reference with regards to all TimeFlows in that container.

container	Reference to the container. Type: Reference to Container Constraints: Required: true
location	A semantic description of a TimeFlow location. Type: string Constraints: Default: LATEST_POINT Acceptable values: LATEST_POINT, LATEST_SNAPSHOT Create: optional Update: optional

TimeflowPointLocation

TimeFlow point based on a database-specific identifier (SCN, LSN, etc).

location	The TimeFlow location. Type: string Constraints: Required: true
timeflow	Reference to TimeFlow containing this location. Type: Reference to Timeflow Constraints: Required: true

TimeflowPointBookmark

bookmark	Reference to the bookmark. Type: Reference to TimeflowBookmark Constraints: Required: true
----------	--

TimeflowPointBookmarkTag

container	Reference to the container. Type: Reference to Container Constraints: Required: true
tag	The name of the tag. Type: string Constraints: Required: true

Timeflow API Objects: Timeflow, snapshot, timeflowRanges

The sample code provides information on the Timeflow, timeflowRanges and snapshot objects for the respective VDB.

Filename: *flows.sh*

```
$ ./flows.sh VBITT10
Session API

Login API
Login {"type":"OKResult","status":"OK","result":"USER-2","job":null,"action":null}
Source: VBITT10
container reference: ORACLE_DB_CONTAINER-75

Timeflows API
timeflow names:
DB_PROVISION@2016-09-27T13:15:18
DB_ROLLBACK@2016-09-28T00:33:54
DB_ROLLBACK@2016-09-28T00:38:33
DB_ROLLBACK@2016-09-28T00:41:44
```

Select Timeflow Name (copy-n-paste from the above list):

```
DB_ROLLBACK@2016-09-28T00:33:54
timeflow reference: ORACLE_TIMEFLOW-97
```

TimeflowRanges for this Timeflow ...

```
{
  "type": "ListResult",
  "status": "OK",
  "result": [
    {
      "type": "TimeflowRange",
      "startPoint": {
        "type": "OracleTimeflowPoint",
        "location": "5475918",
        "timestamp": "2016-09-28T04:35:39.000Z",
        "timeflow": "ORACLE_TIMEFLOW-97"
      },
      "endPoint": {
        "type": "OracleTimeflowPoint",
        "location": "5476181",
        "timestamp": "2016-09-28T04:37:53.000Z",
        "timeflow": "ORACLE_TIMEFLOW-97"
      },
      "provisionable": true
    }
  ],
}
```

```

"job":null,
"action":null,
"total": 1,
"overflow": false
}

```

Snapshot per Timeflow ...

snapshots:

@2016-09-28T04:35:39.826Z

@2016-09-28T04:37:38.537Z

@2016-09-28T04:37:51.273Z

Select Snapshot Name (copy-n-paste from the above list):

```

@2016-09-28T04:37:38.537Z
snapshot reference: ORACLE_SNAPSHOT-152
{
  "type": "OracleSnapshot",
  "reference": "ORACLE_SNAPSHOT-152",
  "namespace":null,
  "name": "@2016-09-28T04:37:38.537Z",
  "consistency": "CRASH_CONSISTENT",
  "missingNonLoggedData": false,
  "container": "ORACLE_DB_CONTAINER-75",
  "creationTime": "2016-09-28T04:37:38.537Z",
  "firstChangePoint": {
    "type": "OracleTimeflowPoint",
    "location": "5476152",
    "timestamp":null,
    "timeflow": "ORACLE_TIMEFLOW-97"
  },
  "latestChangePoint": {
    "type": "OracleTimeflowPoint",
    "location": "5476157",
    "timestamp": "2016-09-28T04:37:38.000Z",
    "timeflow": "ORACLE_TIMEFLOW-97"
  },
  "retention": 0,
  "timeflow": "ORACLE_TIMEFLOW-97",
  "timezone": "US/Eastern,EDT-0400",
  "version": "11.2.0.4.0",
  "runtime": {
    "type": "OracleSnapshotRuntime",
    "provisionable": true,
    "missingLogs":null
  },
  "temporary": false,
  "fromPhysicalStandbyVdb": false,
  "fractionTimeflows":null,
  "redoLogSizeInBytes": 52428800
}

```

```
Done
```

The following scripts demonstrate REFRESH and RESET/ROLLBACK for the Timeflow types of timestamp, snapshot and scn/lsn.

Filename: `vdb_refresh_timestamp.sh`

Filename: `vdb_refresh_snapshot.sh`

Filename: `vdb_refresh_scn.sh`

Filename: `vdb_rollback_timestamp.sh`

Filename: `vdb_rollback_snapshot.sh`

Filename: `vdb_rollback_scn.sh`

Usage: `./vdb_[all_the_above_scripts].sh [source_vdb_name]`

Sample Usage: Copy-and-Paste the timeflow name to the "Select timeflow Name:" prompt. Enter any timestamp between the startPoint and endPoint values using the format:

`[yyyy] - [MM] - [dd]T[HH] : [mm] : [ss] . [SSS]Z`

```
$ ./vdb_rollback_timestamp.sh VBITT
Session and Login Successful ...
database container reference: ORACLE_DB_CONTAINER-73

Timeflows API
Timeflow Names:
DB_REFRESH@2016-09-28T00:13:13
DB_PROVISION@2016-09-26T07:00:08
DB_ROLLBACK@2016-09-26T07:45:30
DB_REFRESH@2016-09-26T12:20:57
DB_ROLLBACK@2016-09-27T12:29:47
```

Select Timeflow Name (copy-n-paste from above list):

```
DB_ROLLBACK@2016-09-26T07:45:30

Timeflow Reference: ORACLE_TIMEFLOW-89
TimeflowRanges for this timeflow ...
{
  "type": "ListResult",
  "status": "OK",
  "result": [
    {
      "type": "TimeflowRange",
      "startPoint": {
        "type": "OracleTimeflowPoint",
        "location": "5474012",
        "timestamp": "2016-09-26T15:30:35.000Z",
        "timeflow": "ORACLE_TIMEFLOW-89"
      },
      "endPoint": {
        "type": "OracleTimeflowPoint",
```

```

        "location": "5482482",
        "timestamp": "2016-09-26T21:15:17.000Z",
        "timeflow": "ORACLE_TIMEFLOW-89"
    },
    "provisionable": true
}
],
"job": null,
"action": null,
"total": 1,
"overflow": false
}

```

Timestamp Format "[yyyy][MM][dd]T[HH]:[mm]:[ss].[SSS]Z"

Enter Timestamp between Start and End Point values (exclude quotes):

2016-09-26T15:45:00.000Z

```

json> {
  "type": "OracleRollbackParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointTimestamp",
    "timeflow": "ORACLE_TIMEFLOW-89",
    "timestamp": "2016-09-26T15:45:00.000Z"
  },
  "username": ""
}

```

Job: JOB-515

Current status as of Wed Sep 28 00:49:39 EDT 2016 : RUNNING 0% Completed

. . .

Current status as of Wed Sep 28 00:51:10 EDT 2016 : RUNNING 73% Completed

Current status as of Wed Sep 28 00:51:50 EDT 2016 : RUNNING 96% Completed

Job: JOB-515 COMPLETED 100% Completed ...

Done ...

API Cookbook: Common Tasks, Workflows, and Examples

These topics describe approaches to common tasks and workflows using the Delphix Engine API.

This section covers the following topics:

- [API Cookbook: Authentication](#)
- [API Cookbook: Host Environment Details](#)
- [API Cookbook: List Alerts and List Jobs](#)
- [API Cookbook: List dSources and VDBs](#)
- [API Cookbook: List Snapshots](#)
- [API Cookbook: Example Provision Of An Oracle VDB](#)
- [API Cookbook: Refresh VDB](#)
- [API Cookbook: Rewind a VDB](#)
- [API Cookbook: Stop/Start a VDB](#)
- [API Cookbook: Creating a Database Template in Delphix Self-Service](#)
- [API Cookbook: Creating a Container in Delphix Self-Service](#)
- [API Cookbook: Refreshing a Container in Delphix Self-Service](#)
- [API Cookbook: Creating a User in Delphix Self-Service](#)
- [API Cookbook: Creating a Branch in Delphix Self-Service](#)
- [API Cookbook: Create a Bookmark in Delphix Self-Service](#)
- [API Cookbook: Delete a Bookmark in Delphix Self-Service](#)
- [API Cookbook: Get a Bookmark in Delphix Self-Service](#)
- [API Cookbook: Share a Bookmark in Delphix Self-Service](#)
- [API Cookbook: Update a Bookmark in Delphix Self-Service](#)
- [API Cookbook: Delete Delphix Self-Service Container](#)
- [API Cookbook: Delete Delphix Self-Service Template](#)
- [API Cookbook: uploadUpgrade](#)

API Cookbook: Authentication

This API cookbook recipe describes how to create an authenticated session for using the Delphix Sever web services.

Before you can use any Delphix Web Service API's you need to create a session, and then authenticate the session by providing valid Delphix account credentials.

Create Delphix API Session

```
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/session \
-c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 4,
    "micro": 3
  }
}
EOF

Response
{
  "status": "OK",
  "result": {
    "type": "APISession",
    "version": {
      "type": "APIVersion",
      "major": 1,
      "minor": 4,
      "micro": 3
    },
    "locale": "en_US",
    "client": null
  },
  "job": null
}
EOF
```

Once the session has been established, the next step is to authenticate to the server by executing the `LoginRequest` API. Unauthenticated sessions are prohibited from making any API calls other than this login request. The username can be either a system user or domain user, and the backend will authenticate using the appropriate method. This example illustrates logging in via curl using cookies created when the session was established:

```
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/login \
```

```
-b cookies.txt -c cookies2.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "delphix_user",
  "password": "delphix_pass",
  "target": "DOMAIN"
}
EOF
```

The new cookie (cookie2.txt) will need to be used in subsequent API requests. The login API currently only supports authentication by a password. There is no way to authenticate using any shared key or alternate authentication strategy.

 It is generally recommended to set the API session version to the [highest level supported](#) by your Delphix Engine.

API Cookbook: Host Environment Details

This API cookbook recipe describes how to obtain host environment details using the Delphix Engine API.

To obtain details about target host environments, list available `Environment` objects on the system. These environments can represent either a single host, or an Oracle cluster.

List Environment

```
curl -X GET -k http://delphix-server/resources/json/delphix/environment \
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For single-host environments, the reference can be used to get information about the associated host. It is also possible to get information about all hosts (regardless of whether they are in a single-host environment or cluster) by omitting the `environment` query parameter.

List UNIX Environment

```
curl -X GET -k http://services.cloud.skytap.com:23173/resources/json/delphix/host?
environment=UNIX_HOST_ENVIRONMENT-1 \
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For more information about the content of host objects, see the `/api/#Host` reference on your local Delphix Engine. Depending on the type of the host, additional information may be available through the following types:

- `UnixHost`
- `WindowsHost`

API Cookbook: List Alerts and List Jobs

This API cookbook recipe describes how to obtain lists of jobs and alerts using the Delphix Engine API.

The `List Alerts` and `List Jobs` API calls can both accept the `toDate` and `fromDate` query parameters to limit rows returned. These parameters require the date to be expressed in [ISO 8601](#) format.

List Alerts

```
$ curl -X GET -k http://delphix-server/resources/json/delphix/alert \  
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For more information about the structure of an alert object, see the `/api/#Alert` link on your local Delphix Engine.

List Jobs (using fromDate)

```
$ curl -X GET -k http://delphix-server/resources/json/delphix/job? \  
  addEvents=true&fromDate=2012-11-08T00:00:00.0000Z \  
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For more information about the structure of a job object, see the `/api/#Job` link on your local Delphix Engine.

API Cookbook: List dSources and VDBs

This API cookbook recipe describes how to obtain a list of dSources and VDBs using the Delphix Engine API.

To obtain a list of dSources and VDBs, list available `Container` (also known as `database`) objects on the system:

List Databases

```
$ curl -X GET -k http://delphix-server/resources/json/delphix/database \
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For more information on the structure of a database object, see the `/api/#Container` reference on your local Delphix Engine. The following sub-types are available depending on the type of database:

- `OracleDatabaseContainer`
- `MSSqlDatabaseContainer`

Each database has zero or one source associated with it. This source could be a linked source, indicating that the database is a dSource, or it could be a virtual source, indicating that it is a VDB. If there are no sources, it is a detached dSource. The `parentContainer` property indicates the reference to the parent container, also indicating that the database is a VDB. To get runtime information about the source associated with the dSource or VDB, use the `Source` API with a `database` parameter set to the reference of the database in question.

List Sources

```
$ curl -X GET -k http://delphix-server/resources/json/delphix/source?
  database=DB_CONTAINER-13 \
  -b ~/cookies.txt -H "Content-Type: application/json"
```

If the `virtual` flag is true, the source is a VDB, otherwise it is a dSource. For more information about the contents of a source object, see the `/api/#Source` reference on your local Delphix Engine. The following sub-types are available depending on the type of source:

- `OracleSource`
 - `OracleLinkedSource`
 - `OracleVirtualSource`
- `MSSqlSource`
 - `MSSqlLinkedSource`
 - `MSSqlVirtualSource`

API Cookbook: List Snapshots

This API cookbook recipe describes how to obtain a list of available snapshots for a VDB or dSource.

Snapshots represent points in time where a `sync` operation has occurred on either a dSource or VDB.

Provisioning from snapshots is much faster than provisioning between snapshots, as the latter requires replaying LogSync records to arrive at the requested point. Given a reference to a database, the `snapshot` API can be used to retrieve the set of snapshots within the database. See the topic [API Cookbook: List dSources and VDBs](#) for information on how to obtain the database reference.

List Snapshots

```
curl -X GET -k http://services.cloud.skytap.com:23173/resources/json/delphix/
snapshot?database=ORACLE_DB_CONTAINER-15 \
  -b ~/cookies.txt -H "Content-Type: application/json"
```

For more information about the structure of a snapshot object, see the `/api/#TimeflowSnapshot` reference on your local Delphix Engine. Snapshots, while representing the point where provisioning will be most efficient, are not the only provisionable points within a database. To get a list of all provisioning points, use the `timeflowRange` API. This API is based on a Timeflow, which is the representation of one timeline within a database. Currently, all databases have a single Timeflow, though this may change in the future. To query for the ranges for a particular database, you will need to use the `currentTimeflow` member of the target database.

List Timeflow Ranges

```
curl -X POST -k --data @- http://services.cloud.skytap.com:23173/resources/json/
delphix/timeflow/ORACLE_TIMEFLOW-11/timeflowRanges \
  -b ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "TimeflowRangeParameters"
}
EOF
```

API Cookbook: Example Provision Of An Oracle VDB

This API cookbook recipe demonstrates how to provision an Oracle VDB using the Delphix Engine API.

In order to provision an Oracle VDB using the API, you need to provide a set of parameters of type **OracleProvisionParameters** (having already authenticated as per [API Cookbook: Authentication](#)).

There are a number of parameters you will need to know:

- **Group reference** - See the list operation in `"/api#group"` on your Delphix Engine
- **VDB name** - The name you want the new VDB to be called
- **Mount path** - Where to mount datasets on the target host.
- **DB/unique names** - The Oracle DB and unique names, often the same as the VDB name
- **Instance name/number** - The Oracle instance name and number to use (dictated by your environment, but often VDB name and 1)
- **Repository reference** - See the list operation on `"/api#repository"` on your Delphix Engine
- **TimeFlow point** - See [API Cookbook: List Snapshots](#) for more information on finding a TimeFlow point, as well as the reference at `"/api#TimeflowPoint Parameters`

You will need to use the structure of the OracleProvisionParameters object to fill it out, see `"/api/#OracleProvisionParameters"` for details on which fields are mandatory/optional.

Here is a minimal example using curl to communicate with the API, provisioning a VDB called "EGVDB" (authentication omitted)

```
curl -X POST -k --data @- http://delphix1.company.com/resources/json/delphix/
database/provision \
  -b cookies.txt -H "Content-Type: application/json" <<EOF
{
  "container": {
    "group": "GROUP-2",
    "name": "EGVDB",
    "type": "OracleDatabaseContainer"
  },
  "source": {
    "type": "OracleVirtualSource",
    "mountBase": "/mnt/provision",
    "allowAutoVDBRestartOnHostReboot": true
  },
  "sourceConfig": {
    "type": "OracleSIConfig",
    "databaseName": "EGVDB",
    "uniqueName": "EGVDB",
    "repository": "ORACLE_INSTALL-3",
    "instance": {
      "type": "OracleInstance",
      "instanceName": "EGVDB",
      "instanceNumber": 1
    }
  }
},
```

```
"timeflowPointParameters": {  
  "type": "TimeflowPointLocation",  
  "timeflow": "ORACLE_TIMEFLOW-123",  
  "location": "3043123"  
},  
"type": "OracleProvisionParameters"  
}  
EOF
```

API Cookbook: Refresh VDB

This API cookbook recipe describes how to refresh a VDB using the Delphix Engine API.

To refresh a VDB you need a reference to the `Database` object, the location of the point to which you wish to refresh and the reference container associated with the object. See the topic [API Cookbook: List dSources and VDBs](#) for information on how to obtain the database reference and current Timeflow. The Timeflow point can be specified either by timestamp, by location (SCN), semantic location or Timeflow bookmark. The `TimeflowPointSemantic` type allows you to specify a semantically meaningful Timeflow location (i.e. the latest snapshot or the latest Timeflow point). The `TimeflowPointBookmark` type allows you to reference a previously created Timeflow bookmark. See [API Cookbook: List Snapshots](#) topic for information on how to determine provisionable points in the parent database.

Refresh VDB

```
curl -v -X POST -k --data @- http://delphix-server/resources/json/delphix/database/
ORACLE_DB_CONTAINER-13/refresh \
  -b ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "OracleRefreshParameters",
  "timeflowPointParameters": {
    "type": "TimeflowPointSemantic",
    "container": "ORACLE_DB_CONTAINER-1",
    "timeflow": "ORACLE_TIMEFLOW-13",
    "location": "LATEST_SNAPSHOT"
  }
}
EOF
```

For more information about the content of refresh parameters, see the `/api/#RefreshParameters` reference on your local Delphix Engine. Depending on the type of the database, the following parameter types are available:

- `OracleRefreshParameters`
- `MSSqlRefreshParameters`

API Cookbook: Rewind a VDB

This API cookbook recipe describes how to rewind a VDB using the Delphix Engine API.

To rewind a VDB, you need a reference to the `Database` object. See the topic, [API Cookbook: List dSources and VDBs](#), for information on how to obtain the database reference. The following sample script includes a working example for creating a session, authenticating to the Delphix Engine, and rewinding the VDB. Please update the script variables to match your environment before using it.

```
#!/bin/bash
#
# sample script to start or stop a VDB.
#
# set this to the FQDN or IP address of the Delphix Engine
DE="192.168.2.131"
# set this to the Delphix admin user name
DELPHIX_ADMIN="delphix_admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# set this to the object reference for the VDB
VDB="ORACLE_DB_CONTAINER-57"
#
# create our session
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/session \

  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 4,
    "micro": 3
  }
}
EOF
{
  "status":"OK",
  "result": {
    "type":"APISession",
    "version": {
      "type": "APIVersion",
      "major": 1,
      "minor": 4,
      "micro": 3
    },
    "locale": "en_US",
    "client": null
  },
  "job": null
}
```

```

EOF
echo
#
# authenticate to the DE
$ curl -s -X POST -k --data @- http://delphix-server/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "delphix_username",
  "password": "delphix_password"
}

EOF
echo
#
# rewind VDB
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/database/${VDB}/
rollback \
  -b ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "OracleRollbackParameters",
  "timeflowPointParameters": {
    "type" : "TimeflowPointSnapshot",
    "snapshot" : "ORACLE_SNAPSHOT-172"
  }
}
EOF
echo

```

 While rewinding a VDB, you can use different parameter types. In the above example, "timeflowPointParameters" type is used as "TimeflowPointSnapshot" and an appropriate snapshot name is provided. Instead of "TimeflowPointSnapshot", you can also choose from "TimeflowPointLocation" or "TimeflowPointTimestamp" or "TimeflowPointBookmark" etc. and pass the relevant parameters.

You can list your Snapshots by following the instructions on [API Cookbook: List Snapshots](#).

API Cookbook: Stop/Start a VDB

This API cookbook recipe describes how to stop and start a VDB using the Delphix Engine API.

To stop or start a VDB, you need a reference to the `Database` object. See the topic, [API Cookbook: List dSources and VDBs](#), for information on how to obtain the database reference. The following script example includes working examples for creating a session, authenticating to the Delphix Engine, and stopping or starting a VDB. Please update the script variables to match your environment before using it. This script requires a single argument which is 'start' or 'stop'.

```
#!/bin/bash
#
# sample script to start or stop a VDB.
#
# set this to the FQDN or IP address of the Delphix Engine
DE="192.168.2.131"
# set this to the Delphix admin user name
DELPHIX_ADMIN="delphix_admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# set this to the object reference for the VDB
VDB="ORACLE_VIRTUAL_SOURCE-5"
#
# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 11,
    "micro": 8
  }
}
EOF
echo
#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -c ~/cookies.txt -b ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
  "password": "${DELPHIX_PASS}"
}
EOF
echo
#
# start or stop the vdb based on the argument passed to the script
case $1 in
```

```
start)
  curl -s -X POST -k http://${DE}/resources/json/delphix/source/${VDB}/start \
    -c ~/cookies.txt -b ~/cookies.txt -H "Content-Type: application/json"
;;
stop)
  curl -s -X POST -k http://${DE}/resources/json/delphix/source/${VDB}/stop \
    -c ~/cookies.txt -b ~/cookies.txt -H "Content-Type: application/json"
;;
*)
  echo "Unknown option: $1"
;;
esac
echo
```

API Cookbook: Creating a Database Template in Delphix Self-Service

 Jet Stream is now known as Delphix Self-Service.

Delphix Self-Service administrators can use this API cookbook recipe to create a database template on Delphix Self-Service using the Delphix Engine API.

 **The following script is for educational and demonstration purposes only and is not supported by Delphix.**

The following script can be downloaded by selecting [createDBTemplate.sh](#).

Create Self-Service Database Template

```

1  #!/bin/bash
2
3  # A sample script for calls to the API. This one creates a Jet Stream
4  Template.
5
6  ##### Constants
7
8  # Describes a Delphix software revision.
9  # Please change version are per your Delphix Engine CLI, if different
10 VERSION="1.8.0"
11
12 ##### Default Values. These can be overwritten with optional arguments.
13 engine="172.16.151.154"
14 username="delphix_admin"
15 password="landshark"
16
17 ##examples##
18 # Create template with mandatory params
19 # ./createBranch.sh -d 172.16.151.154 -u delphix_admin:landshark -n
20 <sourceName> <templateName> <containerName>
21 #Ex ./createBranch.sh -d 172.16.151.154 -u delphix_admin:landshark -n
22 oraclesrc template1 ORACLE_DB_CONTAINER-191
23 # Create template with adding optional params, Notes and Description
24 # ./createDBTemplate.sh -n <sourceName> -N "<templateNotes>" -D
25 "<AnyDescription>" <templateName> <containerName>
26
27 ## NOTE: This script is to add one source per template and it will not add
28 any properties for template, container or source.
29
30 ##### Functions
31
32 # Help Menu
33 function usage {
34     echo "Usage: createDBTemplate.sh [[-h] | options...] <template_name>
35     <source_container>"

```

```

30     echo "Create a Jet Stream Dat Template."
31     echo ""
32     echo "Positional arguments"
33     echo "  <template_name>"
34     echo "  <source_container>"
35     echo ""
36     echo "Optional Arguments:"
37     echo "  -h                Show this message and exit"
38     echo "  -d                Delphix engine IP address or host name,
otherwise revert to default"
39     echo "  -u USER:PASSWORD Server user and password, otherwise revert
to default"
40     echo "  -n                source name to display on JS template
screen"
41     echo "  -N                template notes, if any. Type: String"
42     echo "  -D                source description, if any. Type: String"
43 }
44
45 # Create Our Session, including establishing the API version.
46 function create_session
47 {
48     # Pulling the version into parts. The {} are necessary for string
manipulation.
49     # Strip out longest match following "." This leaves only the major
version.
50     major=${VERSION%%.*}
51     # Strip out the shortest match preceding "." This leaves minor.micro.
52     minorMicro=${VERSION#*.}
53     # Strip out the shortest match following "." This leaves the minor
version.
54     minor=${minorMicro%.*}
55     # Strip out the longest match preceding "." This leaves the micro
version.
56     micro=${VERSION##*.}
57
58     # Quick note about the <<- . If the redirection operator << is followed
by a - (dash), all leading TAB from the document data will be
59     # ignored. This is useful to have optical nice code also when using
here-documents. Otherwise you must have the EOF be on a line by itself,
60     # no parens, no tabs or anything.
61
62     echo "creating session..."
63     result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/
json/delphix/session \
64         -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
65     {
66         "type": "APISession",
67         "version": {
68             "type": "APIVersion",
69             "major": $major,
70             "minor": $minor,
71             "micro": $micro
72     }

```

```

73     }
74     EOF)
75
76     check_result
77 }
78
79 # Authenticate the DE for the provided user.
80 function authenticate_de
81 {
82     echo "authenticating delphix engine..."
83     echo ${engine}
84     echo ${username}
85     echo ${password}
86     result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/
87 json/delphix/login \
88     {
89         "type": "LoginRequest",
90         "username": "${username}",
91         "password": "${password}"
92     }
93     EOF)
94
95     check_result
96 }
97
98 function create_template
99 {
100
101     paramString="\type\": \"JSDataTemplateCreateParameters\", \"name\":
102 \"$templateName\",
103
104     if [[ -n $templatetnotes ]]
105     then
106         paramString="$paramString \"notes\": \"$templatetnotes\","
107     fi
108
109     paramString="$paramString \"dataSources\": [{\"type\":
110 \"JSDataSourceCreateParameters\",
111         \"container\": \"$sourceContainer\",
112         \"source\": {\"type\": \"JSDataSource\",
113         \"priority\": 1,
114         \"name\": \"$sourcename\""}
115
116     if [[ -n $sourcedesc ]]
117     then
118         paramString="$paramString ,\"description\": \"$sourcedesc\"}]"
119     else
120         paramString="$paramString }]"
121     fi
122
123     result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/
124 delphix/jetstream/template \

```

```

122     -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
123     {
124         $paramString
125     }
126     EOF)
127
128     check_result
129
130     echo "New JetStream template $templateName successfully created"
131
132 }
133
134 # Check the result of the curl. If there are problems, inform the user
135 # then exit.
136 function check_result
137 {
138     exitStatus=$?
139     if [ $exitStatus -ne 0 ]
140     then
141         echo "command failed with exit status $exitStatus"
142         exit 1
143     elif [[ $result != *"OKResult"* ]]
144     then
145         echo ""
146         echo $result
147         exit 1
148     fi
149 }
150 ##### Main
151
152 while getopts "u:d:n:N:D:h" flag; do
153     case "$flag" in
154         u )         username=${OPTARG%:*}
155                   password=${OPTARG##*:}
156                   ;;
157         d )         engine=$OPTARG
158                   ;;
159         n )         sourcename=$OPTARG
160                   ;;
161         N )         templatenotes=$OPTARG
162                   ;;
163         D )         sourcedesc=$OPTARG
164                   ;;
165         h )         usage
166                   exit
167                   ;;
168         * )         usage
169                   exit 1
170
171     esac
172
173 done

```

```
174
175
176 # Shift the parameters so we only have the positional arguments left
177 shift $((OPTIND-1))
178
179 # Check that there are 2 positional arguments
180 if [ $# != 2 ]
181 then
182     echo "usage1"
183     usage
184     exit 1
185 fi
186
187 # Get the two positional arguments
188 templateName=$1
189 shift
190 sourceContainer=$1
191
192 create_session
193 authenticate_de
194 create_template
```

API Cookbook: Creating a Container in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to create a container on Delphix Self-Service (Jet Stream) using the Delphix Engine API.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createContainer.sh](#).

Create Self-Service Container

```

1  #!/bin/bash
2
3  # A sample script for calls to the API. This one creates a Jet Stream
4  container.
5
6  ##### Constants
7
8  # Describes a Delphix software revision.
9  # Please change version are per your Delphix Engine CLI, if different
10 VERSION="1.11.10"
11
12 ##### Default Values. These can be overwritten with optional arguments.
13 engine="10.110.248.170"
14 username="admin"
15 password="delphix"
16
17 ##examples##
18 # Create container from latest point in time
19 #./createContainer.sh -n "testsource" testcont ORACLE_DB_CONTAINER-269
20 JS_DATA_TEMPLATE-13
21 # Create container from specific bookmark
22 #./createContainer.sh -n "testsource" -b JS_BOOKMARK-77 testcont
23 ORACLE_DB_CONTAINER-269 JS_DATA_TEMPLATE-13
24 # Create container from specific point in time
25 #./createContainer.sh -n "testsource" -t "2016-08-08T10:00:00.000Z" -B
26 JS_BRANCH-50 testcont ORACLE_DB_CONTAINER-269 JS_DATA_TEMPLATE-13
27
28 #NOTE: this script will add one container and assign one owner for the
29 container.
30
31 ##### Functions
32
33 # Help Menu
34 function usage {
35     echo "Usage: createContainer.sh [[-h] | options...] <containername>
36     <vdb> <template>"
37     echo "Create a Jet Stream Bookmark on the given branch."

```

```

33     echo ""
34     echo "Positional arguments"
35     echo " <name>"
36     echo " <container> format JS_DATA_CONTAINER-<n>"
37     echo ""
38     echo "Optional Arguments:"
39     echo " -h                Show this message and exit"
40     echo " -d                Delphix engine IP address or host name,
otherwise revert to default"
41     echo " -u USER:PASSWORD  Server user and password, otherwise revert
to default"
42     echo " -n                SourceName need to display for container.
(Mandatory)"
43     echo " -b                Bookmark name from which need to create
container. If no bookmark is included, the branch will be created at the
latest point in time. Type: string. Format JS_BOOKMARK-<n> (Optional)"
44     echo " -t                The time at which the branch should be
created. This must be accompanied with branch name from which need to pick
up time. Type: date, must be in ISO 8601 extended format [yyyy]-[MM]-
[dd]T[HH]:[mm]:[ss].[SSS]Z"
45     echo " -B                Branch name from which need to create new
container, at specific time. Type: string. Format JS_BRANCH-<n>
(Optional)"
46     echo " -N                Optional container notes, if need to add
any. Type: String"
47     echo " -o                Optional owner, to whom we need to assign
this container. Type: String. Format USER-<n>"
48 }
49
50 # Create Our Session, including establishing the API version.
51 function create_session
52 {
53     # Pulling the version into parts. The {} are necessary for string
manipulation.
54     # Strip out longest match following "." This leaves only the major
version.
55     major=${VERSION%%.*}
56     # Strip out the shortest match preceding "." This leaves minor.micro.
57     minorMicro=${VERSION##*.}
58     # Strip out the shortest match followint "." This leaves the minor
version.
59     minor=${minorMicro%.*}
60     # Strip out the longest match preceding "." This leaves the micro
version.
61     micro=${VERSION###*.}
62
63     # Quick note about the <<-. If the redirection operator << is followed
by a - (dash), all leading TAB from the document data will be
64     # ignored. This is useful to have optical nice code also when using
here-documents. Otherwise you must have the EOF be on a line by itself,
65     # no parens, no tabs or anything.
66
67     echo "creating session..."

```

```

68     result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/
69     json/delphix/session \
70     -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
71     {
72         "type": "APISession",
73         "version": {
74             "type": "APIVersion",
75             "major": $major,
76             "minor": $minor,
77             "micro": $micro
78         }
79     }
80     EOF)
81     check_result
82 }
83
84 # Authenticate the DE for the provided user.
85 function authenticate_de
86 {
87     echo "authenticating delphix engine..."
88     result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/
89     json/delphix/login \
90     -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/
91     json" <<-EOF
92     {
93         "type": "LoginRequest",
94         "username": "${username}",
95         "password": "${password}"
96     }
97     EOF)
98     check_result
99 }
100 function create_container
101 {
102
103     # If there is not timeInput and no bookmark name, we need to use
104     JSTimelinePointLatestTimeInput.
105     if [[ -z $inputTime && -z $bookmark ]]
106     then
107         pointParams="\timelinePointParameters\":{\"
108             \"sourceDataLayout\": \"${template}\",
109             \"type\": \"JSTimelinePointLatestTimeInput\"}"
110
111     # If there is a timeInput and no bookmark name, we need to use Input
112     Time.
113     elif [[ -n $inputTime && -n $branchRef && -z $bookmark ]]
114     then
115         pointParams="\timelinePointParameters\":{\"
116             \"time\": \"${inputTime}\",

```

```

116         \"branch\": \"${branchRef}\",
117         \"type\": \"JSTimelinePointTimeInput\"}
118
119     # If there is a bookmark name and no time input, we need to use
bookmark
120
121     elif [[ -z $inputTime && -n $bookmark ]]
122     then
123         pointParams=\"\"timelinePointParameters\":{
124             \"bookmark\": \"${bookmark}\",
125             \"type\": \"JSTimelinePointBookmarkInput\"}
126
127     fi
128
129     # These are the required parameters.
130
131     paramString=\"\"type\": \"JSDataContainerCreateWithRefreshParameters\",
132                 \"name\": \"${containerName}\",
133                 \"template\": \"${template}\",\"
134
135
136     paramString=\"$paramString \"dataSources\": [{\"type\":
\"JSDataSourceCreateParameters\",
137                 \"container\": \"${VDB}\",
138                 \"source\": {
139                 \"type\": \"JSDataSource\",
140                 \"priority\": 1,
141                 \"name\": \"${sourceName}\"
142
143     if [[ -n $sourcedesc ]]
144     then
145         paramString=\"$paramString ,\"description\": \"${sourcedesc}\"}],\"
146     else
147         paramString=\"$paramString }],\"
148     fi
149
150     if [[ -n $containerNotes ]]
151     then
152         paramString=\"$paramString \"notes\": \"${containerNotes}\",\"
153     fi
154
155     if [[ -n $owners ]]
156     then
157         paramString=\"$paramString \"owners\": [\"${owners}\",\"
158     fi
159
160     paramString=\"$paramString ${pointParams}\"
161
162
163     result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/
delphix/jetstream/container \
164         -b ~/cookies.txt -H \"Content-Type: application/json\" <<-EOF
165     {

```

```

166     $paramString
167 }
168 EOF)
169
170
171 check_result
172
173
174 echo "confirming job completed successfully..."
175 # Get everything in the result that comes after job.
176 temp=${result#*"job\":"}
177 # Get rid of everything after
178 jobRef=${temp%%\}*}
179
180
181 result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/
job/${jobRef} \
182 -b ~/cookies.txt -H "Content-Type: application/json")
183
184 # Get everything in the result that comes after job.
185 temp=${result#*"jobState\":"}
186 # Get rid of everything after
187 jobState=${temp%%\}*}
188
189
190 check_result
191
192 while [ $jobState = "RUNNING" ]
193 do
194     sleep 1
195     result=$(curl -s -X GET -k http://${engine}/resources/json/
delphix/job/${jobRef} \
196 -b ~/cookies.txt -H "Content-Type: application/json")
197
198     # Get everything in the result that comes after job.
199     temp=${result#*"jobState\":"}
200     # Get rid of everything after
201     jobState=${temp%%\}*}
202
203     check_result
204
205 done
206
207 if [ $jobState = "COMPLETED" ]
208 then
209     echo "successfully created container $containerName"
210 else
211     echo "unable to create container"
212     echo result
213 fi
214
215 }
216

```

```

217 # Check the result of the curl. If there are problems, inform the user
218 # then exit.
219 function check_result
220 {
221     exitStatus=$?
222     if [ $exitStatus -ne 0 ]
223     then
224         echo "command failed with exit status $exitStatus"
225         exit 1
226     elif [[ $result != *"OKResult"* ]]
227     then
228         echo ""
229         echo $result
230         exit 1
231     fi
232 }
233 ##### Main
234
235 while getopts "u:d:b:t:B:D:n:N:o:h" flag; do
236     case "$flag" in
237         u )           username=${OPTARG%:*}
238                     password=${OPTARG##*:}
239                     ;;
240         d )           engine=$OPTARG
241                     ;;
242         b )           bookmark=$OPTARG
243                     ;;
244         t )           inputTime=$OPTARG
245                     ;;
246         B )           branchRef=$OPTARG
247                     ;;
248         D )           sourcedesc=$OPTARG
249                     ;;
250         n )           sourceName=$OPTARG
251                     ;;
252         N )           containerNotes=$OPTARG
253                     ;;
254         o )           owners=$OPTARG
255                     ;;
256         h )           usage
257                     exit
258                     ;;
259         * )           usage
260                     exit 1
261     esac
262
263 done
264
265
266 # Shift the parameters so we only have the positional arguments left
267 shift $((OPTIND-1))
268

```

```
269 # Check that there are 3 positional arguments
270 if [ $# != 3 ]
271 then
272     usage
273     exit 1
274 fi
275
276 # Get the three positional arguments
277 containerName=$1
278 shift
279 VDB=$1
280 shift
281 template=$1
282
283 create_session
284 authenticate_de
285 create_container
```

API Cookbook: Refreshing a Container in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to refresh a container in Delphix Self-Service (Jet Stream) using the Delphix Engine API.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [refreshContainer.sh](#).

Refreshing a Self-Service container

```
#!/bin/bash

# A sample script for calls to the API. This one refresh Jet Stream container.

##### Constants

# Describes a Delphix software revision.
# Please change version are per your Delphix Engine CLI, if different
VERSION="1.8.0"

##### Default Values. These can be overwritten with optional arguments.
engine="172.16.151.154"
username="delphix_admin"
password="landshark"

##examples##
# Refresh container from latest point in time of Template
#./refreshContainer.sh -T JS_DATA_TEMPLATE-13 JS_DATA_CONTAINER-20
# Refresh container from specific bookmark
#./refreshContainer.sh -b JS_BOOKMARK-76 JS_DATA_CONTAINER-20
# Refresh container from specific point in time of branch
#./refreshContainer.sh -t "2016-08-08T10:00:00.000Z" -B JS_BRANCH-50
JS_DATA_CONTAINER-20

##### Functions

# Help Menu
function usage {
    echo "Usage: refreshContainer.sh [[-h] | options...] <containername> <template>"
    echo "Create a Jet Stream Bookmark on the given branch."
    echo ""
    echo "Positional arguments"
    echo "  <containerName>"
    echo "  <template>"
    echo ""
    echo "Optional Arguments:"
```

```

    echo " -h                Show this message and exit"
    echo " -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo " -u USER:PASSWORD Server user and password, otherwise revert to default"
    echo " -T                template reference from which need to refresh from
latest point in time"
    echo " -b                Bookmark name from which need to refresh container. If
no bookmark is included, the branch will be created at the latest point in time.
Type: string. Format JS_BOOKMARK-<n> (Optional)"
    echo " -t                The time from where the container should be refreshed.
This must be accompanied with branch name from which need to pick up time. Type:
date, must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
    echo " -B                Branch name from which need to refresh container, at
specific time. Type: string. Format JS_BRANCH-<n> (Optional)"
}

# Create Our Session, including establishing the API version.
function create_session
{
    # Pulling the version into parts. The {} are necessary for string manipulation.
    # Strip out longest match following "." This leaves only the major version.
    major=${VERSION%%.*}
    # Strip out the shortest match preceding "." This leaves minor.micro.
    minorMicro=${VERSION#*.}
    # Strip out the shortest match following "." This leaves the minor version.
    minor=${minorMicro%.*}
    # Strip out the longest match preceding "." This leaves the micro version.
    micro=${VERSION###*.}

    # Quick note about the <<- . If the redirection operator << is followed by a -
(dash), all leading TAB from the document data will be
    # ignored. This is useful to have optical nice code also when using here-
documents. Otherwise you must have the EOF be on a line by itself,
    # no parens, no tabs or anything.

    echo "creating session..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
session \
    -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "APISession",
        "version": {
            "type": "APIVersion",
            "major": $major,
            "minor": $minor,
            "micro": $micro
        }
    }
}
EOF)

check_result
}

```

```

# Authenticate the DE for the provided user.
function authenticate_de
{
    echo "authenticating delphix engine..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
login \
    -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "LoginRequest",
        "username": "${username}",
        "password": "${password}"
    }
    EOF)

    check_result
}

function restore_container
{
    # If there is not timeInput and no bookmark name, we need to use
    JSTimelinePointLatestTimeInput from template.
    if [[ -n $template && -z $inputTime && -z $bookmark ]]
    then
        pointParams="\\"type\\": \\"JSTimelinePointLatestTimeInput\\",
        \\"sourceDataLayout\\": \\"${template}\\\""

    # If there is a timeInput and no bookmark name, we need to use Input Time.

    elif [[ -n $inputTime && -n $branchRef && -z $bookmark && -z $template ]]
    then
        pointParams="\\"type\\": \\"JSTimelinePointTimeInput\\",
        \\"branch\\": \\"${branchRef}\\",
        \\"time\\": \\"${inputTime}\\\""

    # If there is a bookmark name and no time input, we need to use bookmark

    elif [[ -n $bookmark && -z $template && -z $inputTime ]]
    then
        pointParams="\\"type\\": \\"JSTimelinePointBookmarkInput\\",
        \\"bookmark\\": \\"${bookmark}\\\""

    fi

    echo "pointParams" $pointParams

    result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/delphix/
jetstream/container/${containerRef}/restore \
    -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {

```

```

    $pointParams
}
EOF)

check_result

echo "confirming job completed successfully..."
# Get everything in the result that comes after job.
temp=${result#*"job\":"}
# Get rid of everything after
jobRef=${temp%%\}*}

result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/job/${jobRef}
\
-b ~/cookies.txt -H "Content-Type: application/json")

# Get everything in the result that comes after job.
temp=${result#*"jobState\":"}
# Get rid of everything after
jobState=${temp%%\}*}

check_result

while [ $jobState = "RUNNING" ]
do
    sleep 1
    result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/job/$
{jobRef} \
-b ~/cookies.txt -H "Content-Type: application/json")

    # Get everything in the result that comes after job.
    temp=${result#*"jobState\":"}
    # Get rid of everything after
    jobState=${temp%%\}*}

    check_result

done

if [ $jobState = "COMPLETED" ]
then
    echo "successfully refresh container $containerName"
else
    echo "unable to refresh container"
    echo result
fi
}

```

```

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        exit 1
    elif [[ $result != *"OKResult"* ]]
    then
        echo ""
        echo $result
        exit 1
    fi
}

##### Main

while getopts "u:d:T:b:t:B:h" flag; do
    case "$flag" in
        u )           username=${OPTARG%:*}
                    password=${OPTARG###*:}
                    ;;
        d )           engine=$OPTARG
                    ;;
        T )           template=$OPTARG
                    ;;
        b )           bookmark=$OPTARG
                    ;;
        t )           inputTime=$OPTARG
                    ;;
        B )           branchRef=$OPTARG
                    ;;
        h )           usage
                    exit
                    ;;
        * )           usage
                    exit 1
    esac
done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 1 positional arguments
if [ $# != 1 ]
then
    usage
    exit 1

```

fi

Get the one positional arguments

containerRef=\$1

create_session

authenticate_de

restore_container

API Cookbook: Creating a User in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to create a user on Delphix Self-Service (Jet Stream) using the Delphix Engine API.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createJSUser.sh](#).

Creating a Self-Service User

```
#!/bin/bash

# A sample script for calls to the API. This one creates a Jet Stream user.

##### Constants

# Describes a Delphix software revision.
# Please change version are per your Delphix Engine CLI, if different
VERSION="1.11.9"

##### Default Values. These can be overwritten with optional arguments.
engine="10.43.90.86"
username="admin"
password="delphix"

##examples##
# Create user with NATIVE authentication
#./createJSUser.sh -P <password> NATIVE <username>
# Create user with LDAP authentication
#./createJSUser.sh -r <principal> <LDAP username>

##### Functions

# Help Menu
function usage {
echo "Usage: createJSUser.sh [[-h] | options...] <auth> <newjsuser>"
echo "Create a Jet Stream Only user."
echo ""
echo "Positional arguments"
echo " <auth>"
echo " <newjsuser>"
echo ""
echo "Optional Arguments:"
echo " -h Show this message and exit"
echo " -d Delphix engine IP address or host name, otherwise revert to default"
echo " -u USER:PASSWORD Server user and password, otherwise revert to default"
```

```

echo " -P password for NATIVE authentication"
echo " -f firstName of user"
echo " -l lastName of user"
echo " -e emailAddress of user"
echo " -o homePhoneNumber of user"
echo " -m mobilePhoneNumber of user"
echo " -w workPhoneNumber of user"
echo " -r principal for LDAP authentication"
}

# Create Our Session, including establishing the API version.
function create_session
{
# Pulling the version into parts. The {} are necessary for string manipulation.
# Strip out longest match following "." This leaves only the major version.
major=${VERSION%%.*}
# Strip out the shortest match preceding "." This leaves minor.micro.
minorMicro=${VERSION#*.}
# Strip out the shortest match following "." This leaves the minor version.
minor=${minorMicro%.*}
# Strip out the longest match preceding "." This leaves the micro version.
micro=${VERSION##*.}

# Quick note about the <<-. If the redirection operator << is followed by a - (dash),
all leading TAB from the document data will be
# ignored. This is useful to have optical nice code also when using here-documents.
Otherwise you must have the EOF be on a line by itself,
# no parens, no tabs or anything.

echo "creating session..."
result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
session \
-c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
{
"type": "APISession",
"version": {
"type": "APIVersion",
"major": $major,
"minor": $minor,
"micro": $micro
}
}
EOF)

check_result
}

# Authenticate the DE for the provided user.
function authenticate_de
{
echo "authenticating delphix engine..."

```

```

result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
login \
-b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
{
"type": "LoginRequest",
"username": "${username}",
"password": "${password}"
}
EOF)

check_result
}

function create_user
{
# Check on authorization type

if [[ $authtype = "NATIVE" && -n $userpwd ]]
then
pointParams="\authenticationType\": \"$authtype\",
\"credential\": {
\"type\": \"PasswordCredential\",
\"password\": \"$userpwd\""

elif [[ $authtype = "LDAP" && -n $principal ]]
then
pointParams="\authenticationType\": \"$authtype\",
\"principal\": \"$principal\""

fi

# These are the required parameters.
paramString="
\"type\": \"User\",
\"name\": \"${newjsuser}\"

# Fill in optional parameters if there are any.
if [[ -n $firstname ]]
then
paramString="$paramString \"firstName\": \"$firstname\","
fi

if [[ -n $lastname ]]
then
paramString="$paramString \"lastName\": \"$lastname\","
fi

if [[ -n $emailaddress ]]
then
paramString="$paramString \"emailAddress\": \"$emailaddress\","
fi

```

```

if [[ -n $homephone ]]
then
paramString="$paramString \"homePhoneNumber\": \"$homephone\","
fi

if [[ -n $mobilephone ]]
then
paramString="$paramString \"mobilePhoneNumber\": \"$mobilephone\","
fi

if [[ -n $workphone ]]
then
paramString="$paramString \"workPhoneNumber\": \"$workphone\","
fi

paramString="$paramString
${pointParams}"

result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/delphix/user \
-b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
{
$paramString
}
EOF)

check_result

# Extracting USER ID from result
temp=${result#*\ "result\":"\}
userRef=${temp%%\ "*}

echo "New user $newjsuser successfully created"

##### ROLE-3 is Jet Stream Role

result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/delphix/
authorization \
-b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
{
"type": "Authorization",
"role": "ROLE-3",
"target": "$userRef",
"user": "$userRef"
}
EOF)

check_result

echo "Assigned Jet Stream Role to user $newjsuser"

```

```

}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
exitStatus=$?
if [ $exitStatus -ne 0 ]
then
echo "command failed with exit status $exitStatus"
exit 1
elif [[ $result != *"OKResult"* ]]
then
echo ""
echo $result
exit 1
fi
}

##### Main

while getopts "u:d:P:r:f:l:e:o:m:w:h" flag; do
case "$flag" in
u ) username=${OPTARG%:*}
password=${OPTARG##*:}
;;
d ) engine=$OPTARG
;;
P ) userpwd=$OPTARG
;;
r ) principal=$OPTARG
;;
f ) firstname=$OPTARG
;;
l ) lastname=$OPTARG
;;
e ) emailaddress=$OPTARG
;;
o ) homephone=$OPTARG
;;
m ) mobilephone=$OPTARG
;;
w ) workphone=$OPTARG
;;
h ) usage
exit
;;
* ) usage
exit 1
esac
done

```

```
# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 2 positional arguments
if [ $# != 2 ]
then
echo "usage1"
usage
exit 1
fi

# Get the two positional arguments
authtype=$1
shift
newjsuser=$1

create_session
authenticate_de
create_user
```

API Cookbook: Creating a Branch in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to create a branch on Delphix Self-Service (Jet Stream) using the Delphix Engine API.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [createBranch.sh](#).

Create Self-Service Branch

```

1  #!/bin/bash
2
3  # A sample script for calls to the API. This one creates a Jet Stream
4  branch.
5
6  ##### Constants
7
8  # Describes a Delphix software revision.
9  # Please change version are per your Delphix Engine CLI, if different
10 VERSION="1.8.0"
11
12 ##### Default Values. These can be overwritten with optional arguments.
13 engine="172.16.151.154"
14 username="delphix_admin"
15 password="landshark"
16
17 ##examples##
18 # Create branch from latest point in time
19 #./createBranch.sh -d 172.16.151.154 -u delphix_admin:landshark newbranch
20 JS_DATA_CONTAINER-20
21 # Create branch from specific bookmark
22 #./createBranch.sh -d 172.16.151.154 -u delphix_admin:landshark -b
23 bookmarkname newbranch JS_DATA_CONTAINER-20
24 # Create branch from specific point in time
25 #./createBranch.sh -d 172.16.151.154 -u delphix_admin:landshark -t
26 "2016-07-27T01:45:56.453Z" -B JS_BRANCH-7 newbranch JS_DATA_CONTAINER-20
27
28 ##### Functions
29
30 # Help Menu
31 function usage {
32     echo "Usage: createBranch.sh [[-h] | options...] <name> <container>"
33     echo "Create a Jet Stream Bookmark on the given branch."
34     echo ""
35     echo "Positional arguments"
36     echo "  <name>"
37     echo "  <container> format JS_DATA_CONTAINER-<n>"

```

```

35     echo ""
36     echo "Optional Arguments:"
37     echo " -h                Show this message and exit"
38     echo " -d                Delphix engine IP address or host name,
otherwise revert to default"
39     echo " -u USER:PASSWORD  Server user and password, otherwise revert
to default"
40     echo " -b                Bookmark name from which need to create
branch. If no bookmark is included, the branch will be created at the
latest point in time. Type: string. Format JS_BOOKMARK-<n>"
41     echo " -t                The time at which the branch should be
created. If no time is included, the branch will be created at the latest
point in time. Type: date, must be in ISO 8601 extended format [yyyy]-
[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
42     echo " -B                Branch name from which need to create new
branch, at specific time."
43 }
44
45 # Create Our Session, including establishing the API version.
46 function create_session
47 {
48     # Pulling the version into parts. The {} are necessary for string
manipulation.
49     # Strip out longest match following "." This leaves only the major
version.
50     major=${VERSION%%.*}
51     # Strip out the shortest match preceding "." This leaves minor.micro.
52     minorMicro=${VERSION#*.}
53     # Strip out the shortest match followint "." This leaves the minor
version.
54     minor=${minorMicro%.*}
55     # Strip out the longest match preceding "." This leaves the micro
version.
56     micro=${VERSION##*.}
57
58     # Quick note about the <<- . If the redirection operator << is followed
by a - (dash), all leading TAB from the document data will be
59     # ignored. This is useful to have optical nice code also when using
here-documents. Otherwise you must have the EOF be on a line by itself,
60     # no parens, no tabs or anything.
61
62     echo "creating session..."
63     result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/
json/delphix/session \
64         -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
65     {
66         "type": "APISession",
67         "version": {
68             "type": "APIVersion",
69             "major": $major,
70             "minor": $minor,
71             "micro": $micro
72     }

```

```

73     }
74     EOF)
75
76     check_result
77 }
78
79 # Authenticate the DE for the provided user.
80 function authenticate_de
81 {
82     echo "authenticating delphix engine..."
83     result=$(curl -s -S -X POST -k --data @- http://{engine}/resources/
84 json/delphix/login \
85     -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/
86 json" <<-EOF
87     {
88         "type": "LoginRequest",
89         "username": "${username}",
90         "password": "${password}"
91     }
92     EOF)
93
94     check_result
95 }
96
97 function create_branch
98 {
99     # If there is not timeInput and no bookmark name, we need to use
100 JSTimelinePointLatestTimeInput.
101     if [[ -z $inputTime && -z $bookmark ]]
102     then
103         pointParams="\\"timelinePointParameters\":{
104             \"sourceDataLayout\": \"$container\",
105             \"type\":"JSTimelinePointLatestTimeInput\"}"
106
107     # If there is a timeInput and no bookmark name, we need to use Input
108 Time.
109
110     elif [[ -n $inputTime && -z $bookmark ]]
111     then
112         pointParams="\\"timelinePointParameters\":{
113             \"time\":"$inputTime\",
114             \"branch\":"$branchRef\",
115             \"sourceDataLayout\":"$container\",
116             \"type\":"JSTimelinePointTimeInput\"}"
117
118     # If there is a bookmark name and no time input, we need to use
119 bookmark
120
121     elif [[ -z $inputTime && -n $bookmark ]]
122     then
123         pointParams="\\"timelinePointParameters\":{
124             \"bookmark\":"$bookmark\",

```

```

121         \"type\": \"JSTimelinePointBookmarkInput\"}"}
122
123     fi
124
125     # These are the required parameters.
126     paramString="
127         \"dataContainer\": \"${container}\",
128         \"name\": \"${branchName}\",
129
130
131     paramString="$paramString
132         ${pointParams},
133         \"type\": \"JSBranchCreateParameters\""
134
135
136     result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/
delphix/jetstream/branch \
137         -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
138     {
139         $paramString
140     }
141     EOF)
142
143
144     check_result
145
146
147
148     echo "confirming job completed successfully..."
149     # Get everything in the result that comes after job.
150     temp=${result#*"job\": \""}
151     # Get rid of everything after
152     jobRef=${temp%%\}*}
153
154
155     result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/
job/${jobRef} \
156     -b ~/cookies.txt -H "Content-Type: application/json")
157
158     # Get everything in the result that comes after job.
159     temp=${result#*"jobState\": \""}
160     # Get rid of everything after
161     jobState=${temp%%\}*}
162
163
164     check_result
165
166     while [ $jobState = "RUNNING" ]
167     do
168         sleep 1
169         result=$(curl -s -X GET -k http://${engine}/resources/json/
delphix/job/${jobRef} \
170         -b ~/cookies.txt -H "Content-Type: application/json")

```

```

171
172     # Get everything in the result that comes after job.
173     temp=${result#\*"jobState\":"}
174     # Get rid of everything after
175     jobState=${temp%%\*" }
176
177     check_result
178
179     done
180
181     if [ $jobState = "COMPLETED" ]
182     then
183         echo "successfully created branch $branchName"
184     else
185         echo "unable to create branch"
186         echo result
187     fi
188
189 }
190
191 # Check the result of the curl. If there are problems, inform the user
192 # then exit.
193 function check_result
194 {
195     exitStatus=$?
196     if [ $exitStatus -ne 0 ]
197     then
198         echo "command failed with exit status $exitStatus"
199         exit 1
200     elif [[ $result != *"OKResult"* ]]
201     then
202         echo ""
203         echo $result
204         exit 1
205     fi
206 }
207 ##### Main
208
209 while getopts "u:d:b:t:B:h" flag; do
210     case "$flag" in
211         u )         username=${OPTARG:*}
212                   password=${OPTARG##*:}
213                   ;;
214         d )         engine=$OPTARG
215                   ;;
216         b )         bookmark=$OPTARG
217                   ;;
218         t )         inputTime=$OPTARG
219                   ;;
220         B )         branchRef=$OPTARG
221                   ;;
222         h )         usage

```

```
223             exit
224             ;;
225         * )      usage
226             exit 1
227     esac
228
229 done
230
231
232 # Shift the parameters so we only have the positional arguments left
233 shift $((OPTIND-1))
234
235 # Check that there are 2 positional arguments
236 if [ $# != 2 ]
237 then
238     usage
239     exit 1
240 fi
241
242 # Get the two positional arguments
243 branchName=$1
244 shift
245 container=$1
246
247 create_session
248 authenticate_de
249 create_branch
```

API Cookbook: Create a Bookmark in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to create a bookmark on Delphix Self-Service (Jet Stream) using the Delphix Engine API.

i **The following script is for educational and demonstration purposes only and is not supported by Delphix.**

This script can be downloaded by selecting **createBookmark.sh**.

Creating a bookmark in Self-Service

```
#!/bin/bash

# A sample script for calls to the API. This one creates a Jet Stream bookmark.

##### Constants

# Describes a Delphix software revision.
# Please change version are per your Delphix Engine CLI, if different
VERSION="1.8.0"

##### Default Values. These can be overwritten with optional arguments.
engine="172.16.151.154"
username="delphix_admin"
password="landshark"

shared=false

##example##
./createBookmark.sh -d 172.16.151.154 -u delphix_admin:landshark -p "bookmark test"
-e "2016-07-27T23:38:56.453Z" -t "2016-07-27T01:45:56.453Z" -T
"[tag1,tag2,tag3,tag4,tag5]" bkmrk3 JS_BRANCH-41

##### Functions

# Help Menu
function usage {
    echo "Usage: createBookmark.sh [[-h] | options...] <name> <branch>"
    echo "Create a Jet Stream Bookmark on the given branch."
    echo ""
    echo "Positional arguments"
    echo "  <name>"
    echo "  <branch>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h          Show this message and exit"
```

```

    echo " -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo " -u USER:PASSWORD Server user and password, otherwise revert to default"
    echo " -D                Description of this bookmark. Type: string"
    echo " -e                A policy will automatically delete this bookmark at
this time. If not present the bookmark will be kept until manually deleted. Type:
date, must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
    echo " -s                Present if need to make bookmark in shared mode"
    echo " -t                The time at which the bookmark should be created. If no
time is included, the bookmark will be created at the latest point in time. Type:
date, must be in ISO 8601 extended format [yyyy]-[MM]-[dd]T[HH]:[mm]:[ss].[SSS]Z"
    echo " -T                A set of user-defined labels for this bookmark. No
spaces allowed. Array of Type: string. In format, [tag1,tag2,..] "
}

# Create Our Session, including establishing the API version.
function create_session
{
    # Pulling the version into parts. The {} are necessary for string manipulation.
    # Strip out longest match following "." This leaves only the major version.
    major=${VERSION%%.*}
    # Strip out the shortest match preceding "." This leaves minor.micro.
    minorMicro=${VERSION#*.}
    # Strip out the shortest match following "." This leaves the minor version.
    minor=${minorMicro%.*}
    # Strip out the longest match preceding "." This leaves the micro version.
    micro=${VERSION###*.}

    # Quick note about the <<-. If the redirection operator << is followed by a -
(dash), all leading TAB from the document data will be
    # ignored. This is useful to have optical nice code also when using here-
documents. Otherwise you must have the EOF be on a line by itself,
    # no parens, no tabs or anything.

    echo "creating session..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
session \
        -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "APISession",
        "version": {
            "type": "APIVersion",
            "major": $major,
            "minor": $minor,
            "micro": $micro
        }
    }
    EOF)
    check_result
}

```

```

# Authenticate the DE for the provided user.
function authenticate_de
{
    echo "authenticating delphix engine..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
login \
    -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "LoginRequest",
        "username": "${username}",
        "password": "${password}"
    }
    EOF)

    check_result
}

# Get the branch info so the bookmark to fill in dataLayout
function get_branch
{
    echo "retrieving branch $branchRef to find Source Data Layout..."
    result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/jetstream/
branch/${branchRef} \
    -b ~/cookies.txt -H "Content-Type: application/json")

    check_result

    # Get everything in the result that comes after dataLayout.
    temp=${result#*"dataLayout\":"}
    # Get rid of everything after creat
    dataLayout=${temp%%\}*}

    echo "temp" $temp

    echo "dataLayout" $dataLayout
}

function create_bookmark
{
    get_branch

    # If there is not creation time, we need to use JSTimelinePointLatestTimeInput.
    if [ -z $creationTime ]
    then
        pointParams="\timelinePointParameters\":{\"
            \"sourceDataLayout\": \"$dataLayout\",
            \"type\": \"JSTimelinePointLatestTimeInput\"}"
    else
        pointParams="\timelinePointParameters\":{\"
            \"sourceDataLayout\": \"$dataLayout\",
            \"time\": \"$creationTime\",

```

```

        \ "branch\":\ "${branchRef}\",
        \ "type\":\ "JSTimelinePointTimeInput\"}"
    fi

# These are the required parameters.
paramString="
    \ "bookmark\": {
        \ "branch\": \ "${branchRef}\",
        \ "name\": \ "${bookmarkName}\",
    }

# Fill in optional parameters if there are any.
if [[ -n $description ]]
then
    paramString="$paramString \ "description\": \ "${description}\",
fi

if [[ -n $expiration ]]
then
    paramString="$paramString \ "expiration\": \ "${expiration}\",
fi

if [[ -n $shared ]]
then
    paramString="$paramString \ "shared\": $shared,"
fi

if [[ -n $tags ]]
then
    # Add quotes back to the passed in tags so they are processed correctly.
    tags=${tags//[/[\\]}
    tags=${tags//,/\\,\\}
    tags=${tags//]/\\]}

    paramString="$paramString \ "tags\": $tags,"
fi

paramString="$paramString \ "type\": \ "JSBookmark\"
    },
    ${pointParams},
    \ "type\": \ "JSBookmarkCreateParameters\""}

result=$(curl -s -X POST -k --data @- http://${engine}/resources/json/delphix/
jetstream/bookmark \
    -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
{
    $paramString
}
EOF)

check_result

echo "confirming job completed successfully..."

```

```

# Get everything in the result that comes after job.
temp=${result#*"job\":"}
# Get rid of everything after
jobRef=${temp%%\}*}

result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/job/${jobRef}
\
-b ~/cookies.txt -H "Content-Type: application/json")

# Get everything in the result that comes after job.
temp=${result#*"jobState\":"}
# Get rid of everything after
jobState=${temp%%\}*}

check_result

while [ $jobState = "RUNNING" ]
do
    sleep 1
    result=$(curl -s -X GET -k http://${engine}/resources/json/delphix/job/$
{jobRef} \
-b ~/cookies.txt -H "Content-Type: application/json")

    # Get everything in the result that comes after job.
    temp=${result#*"jobState\":"}
    # Get rid of everything after
    jobState=${temp%%\}*}

    check_result

done

if [ $jobState = "COMPLETED" ]
then
    echo "successfully created bookmark $bookmarkName"
else
    echo "unable to create bookmark"
    echo result
fi
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        exit 1
    elif [[ $result != *"OKResult"* ]]

```

```

    then
        echo ""
        echo $result
        exit 1
    fi
}

##### Main

while getopts "u:d:D:e:s:t:T:h" flag; do
    case "$flag" in
        u )          username=${OPTARG%:*}
                    password=${OPTARG##*:}
                    ;;
        d )          engine=$OPTARG
                    ;;
        D )          description=$OPTARG
                    ;;
        e )          expiration=$OPTARG
                    ;;
        s )          shared=true
                    ;;
        t )          creationTime=$OPTARG
                    ;;
        T )          tags=$OPTARG
                    ;;
        h )          usage
                    exit
                    ;;
        * )          usage
                    exit 1
    esac

    echo "OPTARG" $OPTARG #####

done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there are 2 positional arguments
if [ $# != 2 ]
then
    usage
    exit 1
fi

# Get the two positional arguments
bookmarkName=$1
shift
branchRef=$1

```

```
create_session  
authenticate_de  
create_bookmark
```

API Cookbook: Delete a Bookmark in Delphix Self-Service

This API cookbook recipe describes how to delete a bookmark in Delphix Self-Service (Jet Stream).

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

Deleting a bookmark in Self-Service

```
#!/bin/bash
#
# sample script to delete a bookmark on a Jet Stream container.
#
# Please set the following variables to suit your purposes.
# set this to the FQDN or IP address of the Delphix Engine
DE="ars-dlpx-6010-3.dlpxdc.co"
# set this to the Delphix admin user name
DELPHIX_ADMIN="admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# reference of bookmark you want to delete
BOOKMARK_REF="JS_BOOKMARK-2"
#
# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 6,
    "micro": 2
  }
}
EOF
echo
#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
  "password": "${DELPHIX_PASS}"
}
EOF
echo
#
```

```
# delete the bookmark
curl -s -X DELETE -k http://${DE}/resources/json/delphix/jetstream/bookmark/${BOOKMARK_REF} \
  -b ~/cookies.txt -H "Content-Type: application/json"
echo
```

API Cookbook: Get a Bookmark in Delphix Self-Service

This API cookbook recipe describes how to [get a bookmark](#) in Delphix Self-Service (Jet Stream).

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

```
#!/bin/bash
#
# sample script to get a bookmark on a Jet Stream container.
#
# Please set the following variables to suit your purposes.
# set this to the FQDN or IP address of the Delphix Engine
DE="110.110.200.107"
# set this to the Delphix admin user name
DELPHIX_ADMIN="admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# reference of bookmark you want to get
BOOKMARK_REF="JS_BOOKMARK-2"
#
# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 6,
    "micro": 2
  }
}
EOF
echo
#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
  "password": "${DELPHIX_PASS}"
}
EOF
echo
#
# get the bookmark
```

```
curl -s -X GET -k http://${DE}/resources/json/delphix/jetstream/bookmark/${BOOKMARK_REF} \  
  -b ~/cookies.txt -H "Content-Type: application/json" \  
echo
```

API Cookbook: Share a Bookmark in Delphix Self-Service

Delphix Self-Service administrators can use this API cookbook recipe to share a bookmark in Delphix Self-Service (Jet Stream) using the Delphix Engine API.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [shareBookmark.sh](#).

```
#!/bin/bash

# A sample script for calls to the API. This one shares Bookmark across containers in
# same template.

##### Constants

# Describes a Delphix software revision.
# Please change version are per your Delphix Engine CLI, if different.
VERSION="1.11.10"

##### Default Values. These can be overwritten with optional arguments.
engine="ars-dlpx-6010-3.dlpxdc.co"
username="admin"
password="delphix"

##examples##
# Share Bookmark
#./shareBookmark.sh -a share JS_BOOKMARK-75
# Unshare Bookmark
#./shareBookmark.sh -a unshare JS_BOOKMARK-75

##### Functions

# Help Menu
function usage {
    echo "Usage: shareBookmark.sh [[-h] | options...] <bookmarkName>"
    echo "Share/Unshare JetStream bookmark"
    echo ""
    echo "Positional arguments"
    echo "bookmarkName. Format: JS_BOOKMARK-<n>"
    echo ""
    echo "Optional Arguments:"
    echo "  -h                Show this message and exit"
    echo "  -d                Delphix engine IP address or host name, otherwise
revert to default"
    echo "  -u USER:PASSWORD Server user and password, otherwise revert to default"
    echo "  -a                action to perform on bookmark. Type:String.
Values:share/unshare"
}
}
```

```

# Create Our Session, including establishing the API version.
function create_session
{
    # Pulling the version into parts. The {} are necessary for string manipulation.
    # Strip out longest match following "." This leaves only the major version.
    major=${VERSION%%.*}
    # Strip out the shortest match preceding "." This leaves minor.micro.
    minorMicro=${VERSION#*.}
    # Strip out the shortest match following "." This leaves the minor version.
    minor=${minorMicro%.*}
    # Strip out the longest match preceding "." This leaves the micro version.
    micro=${VERSION###*.}

    # Quick note about the <<-. If the redirection operator << is followed by a -
    (dash), all leading TAB from the document data will be
    # ignored. This is useful to have optical nice code also when using here-
    documents. Otherwise you must have the EOF be on a line by itself,
    # no parens, no tabs or anything.

    echo "creating session..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
session \
    -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "APISession",
        "version": {
            "type": "APIVersion",
            "major": $major,
            "minor": $minor,
            "micro": $micro
        }
    }
    EOF)
    check_result
}

# Authenticate the DE for the provided user.
function authenticate_de
{
    echo "authenticating delphix engine..."
    result=$(curl -s -S -X POST -k --data @- http://${engine}/resources/json/delphix/
login \
    -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<-EOF
    {
        "type": "LoginRequest",
        "username": "${username}",
        "password": "${password}"
    }
    EOF)
    check_result
}

```

```

}

function bookmark_action
{
    # Change share mode of bookmark

    if [[ $action = "share" ]]
    then
        result=$(curl -s -X POST -k --data @- http://{engine}/resources/json/
delphix/jetstream/bookmark/${bookmarkName}/${action} \
        -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
        {}
        EOF)

        check_result

        echo "Bookmark ${bookmarkName} is now in shared mode"

    elif [[ $action = "unshare" ]]
    then
        result=$(curl -s -X POST -k --data @- http://{engine}/resources/json/
delphix/jetstream/bookmark/${bookmarkName}/${action} \
        -b ~/cookies.txt -H "Content-Type: application/json" <<-EOF
        {}
        EOF)

        check_result

        echo "Bookmark ${bookmarkName} is now in not-share mode"

    fi
}

# Check the result of the curl. If there are problems, inform the user then exit.
function check_result
{
    exitStatus=$?
    if [ $exitStatus -ne 0 ]
    then
        echo "command failed with exit status $exitStatus"
        exit 1
    elif [[ $result != *"OKResult"* ]]
    then
        echo ""
        echo $result
        exit 1
    fi
}

##### Main

while getopt u:d:a:h flag; do

```

```
case "$flag" in
  u )      username=${OPTARG%:*}
           password=${OPTARG##*:}
           ;;
  d )      engine=$OPTARG
           ;;
  a )      action=$OPTARG
           ;;
  h )      usage
           exit
           ;;
  * )      usage
           exit 1
esac

done

# Shift the parameters so we only have the positional arguments left
shift $((OPTIND-1))

# Check that there is 1 positional arguments
if [ $# != 1 ]
then
  usage
  exit 1
fi

# Get the one positional arguments
bookmarkName=$1

create_session
authenticate_de
bookmark_action
```

API Cookbook: Update a Bookmark in Delphix Self-Service

This API cookbook recipe describes how to [update a Bookmark](#) in Delphix Self-Service (Jet Stream). Note that the following example includes updating the "tags" on a Delphix Self-Service bookmark.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

```
#!/bin/bash
#
# sample script to update a bookmark on a Jet Stream container.
#
# Please set the following variables to suit your purposes.
# set this to the FQDN or IP address of the Delphix Engine
DE="ars-dlpx-6010-3.dlpxdc.co"
# set this to the Delphix admin user name
DELPHIX_ADMIN="admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# reference of bookmark you want to update
BOOKMARK_REF="JS_BOOKMARK-2"
#
# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 6,
    "micro": 2
  }
}
EOF
echo
#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
  "password": "${DELPHIX_PASS}"
}
EOF
echo
#
# Update the bookmark. Note that only fields you want to change must be included in
the bookmark
```

```
# json.
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/jetstream/bookmark/${BOOKMARK_REF} \
  -b ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "JSBookmark",
  "tags": ["tagA", "tabB"]
}
EOF
echo
```

API Cookbook: Delete Delphix Self-Service Container

This API cookbook recipe describes how to delete a container in Delphix Self-Service.

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix.

This script can be downloaded by selecting [deleteContainer.sh](#)

Delete Delphix Self-Service Container

```
#!/bin/bash
#
# sample script to delete a bookmark on a Jet Stream container.
#
# Please set the following variables to suit your purposes.
# set this to the FQDN or IP address of the Delphix Engine
DE="ars-6010.dlpxdc.co"
# set this to the Delphix admin user name
DELPHIX_ADMIN="admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# reference of container you want to delete
CONTAINER_REF="JS_DATA_CONTAINER-1"
#
# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 6,
    "micro": 2
  }
}
EOF
echo
#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
  "password": "${DELPHIX_PASS}"
```

```
}  
EOF  
echo  
  
#  
# delete the bookmark  
curl -s -X DELETE -k http://${DE}/resources/json/delphix/jetstream/container/${  
{CONTAINER_REF} \  
    -b ~/cookies.txt -H "Content-Type: application/json"  
echo
```

API Cookbook: Delete Delphix Self-Service Template

This API cookbook recipe describes how to delete a template in Delphix Self-Service

ⓘ The following script is for educational and demonstration purposes only and is not supported by Delphix

This script can be downloaded by selecting [deleteTemplate.sh](#).

Delete Delphix Self-Service Template

```
#!/bin/bash
#
# sample script to delete a bookmark on a Jet Stream container.
#
# Please set the following variables to suit your purposes.
# set this to the FQDN or IP address of the Delphix Engine
DE="ars-dlpx-6010.dlpxdc.co"
# set this to the Delphix admin user name
DELPHIX_ADMIN="admin"
# set this to the password for the Delphix admin user
DELPHIX_PASS="delphix"
# reference of template you want to delete
TEMPLATE_REF="JS_DATA_TEMPLATE-1"
#

# create our session
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/session \
  -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 6,
    "micro": 2
  }
}
EOF
echo

#
# authenticate to the DE
curl -s -X POST -k --data @- http://${DE}/resources/json/delphix/login \
  -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "LoginRequest",
  "username": "${DELPHIX_ADMIN}",
```

```
    "password": "${DELPHIX_PASS}"
}
EOF
echo

#
# delete the bookmark
curl -s -X DELETE -k http://${DE}/resources/json/delphix/jetstream/template/${
TEMPLATE_REF} \
    -b ~/cookies.txt -H "Content-Type: application/json"
echo
```

API Cookbook: uploadUpgrade

This API cookbook recipe describes how to use uploadUpgrade.

```
curl -s -X POST -k --data @- http://delphix.engine/resources/json/delphix/session -c
~/cookies.txt -H "Content-Type: application/json" <<EOF
{
  "type": "APISession",
  "version": {
    "type": "APIVersion",
    "major": 1,
    "minor": 10,
    "micro": 0
  }
}
EOF

curl -s -X POST -k --data @- httpk --data @- http://delphix.engine/resources/json/
delphix/login -b ~/cookies.txt -c ~/cookies.txt -H "Content-Type: application/json"
<<EOF
{
  "type": "LoginRequest",
  "username": "sysadmin",
  "password": "sysadmin"
}
EOF

curl -s -X POST -F file=@upgrade_image_path http://delphix.engine/resources/json/
system/uploadUpgrade -b ~/cookies.txt
```

Kerberos APIs

The following topics provide information related to Kerberos APIs.

- [API Cookbook: ASEDBConfig](#)
- [API Cookbook: ASEhostEnvironmentParameters](#)
- [API Cookbook: SAP ASE Instance](#)
- [API Cookbook: ASELinkData](#)
- [API Cookbook: EnvironmentUser](#)
- [API Cookbook: KerberosConfig](#)
- [API Cookbook: KerberosCredential](#)
- [API Cookbook: KerberosKDC](#)

API Cookbook: ASEDBConfig

This API cookbook recipe describes how to configure your SAP ASE database using the Delphix Engine API.

```
{
  "name": "ASEDBConfig",
  "description": "A SAP ASE Database Config.",
  "abstract": true,
  "extends": {
    "$ref": "/delphix-source-config.json"
  },
  "properties": {
    "databaseName": {
      "type": "string",
      "description": "The name of the database.",
      "create": "required",
      "update": "optional",
      "pattern": "^[a-zA-Z0-9_]+$",
      "maxLength": 30
    },
    "user": {
      "type": "string",
      "description": "The username of the database user.",
      "update": "optional",
      "maxLength": 256
    },
    "credentials": {
      "type": "object",
      "description": "The password of the database user.",
      "$ref": "/delphix-credential.json",
      "update": "optional"
    },
    "repository": {
      "type": "string",
      "description": "The object reference of the source repository.",
      "format": "objectReference",
      "referenceTo": "/delphix-ase-instance.json",
      "create": "required",
      "update": "optional"
    }
  }
}
```

API Cookbook: ASEHostEnvironmentParameters

This API cookbook recipe describes how to configure your SAP ASE host environment parameters using the Delphix Engine API.

```
{
  "name": "ASEHostEnvironmentParameters",
  "description": "SAP ASE host environment parameters.",
  "extends": {
    "$ref": "/delphix-typed-object.json"
  },
  "properties": {
    "dbUser": {
      "type": "string",
      "description": "The username of the database user.",
      "create": "optional",
      "update": "optional",
      "maxLength": 256
    },
    "credentials": {
      "type": "object",
      "description": "The credentials of the database user.",
      "$ref": "/delphix-credential.json",
      "create": "required",
      "update": "optional",
      "properties": {
        "type": {
          "type": "string",
          "description": "Object type.",
          "required": true,
          "format": "type",
          "default": "PasswordCredential"
        }
      }
    }
  }
}
```

API Cookbook: SAP ASE Instance

This API cookbook recipe describes how to configure your SAP ASE instance using the Delphix Engine API.

```
{
  "name": "ASEInstance",
  "description": "The SAP ASE source repository.",
  "extends": {
    "$ref": "/delphix-source-repository.json"
  },
  "properties": {
    "instanceName": {
      "type": "string",
      "description": "The name of the SAP ASE instance.",
      "create": "required"
    },
    "installationPath": {
      "type": "string",
      "description": "The SAP ASE instance home.",
      "create": "required",
      "update": "optional"
    },
    "ports": {
      "type": "array",
      "description": "The network ports for connecting to the SAP ASE
instance.",
      "items": {
        type: "integer"
      },
      "create": "required",
      "update": "optional"
    },
    "instanceOwner": {
      "type": "string",
      "description": "The username of the account the SAP ASE instance is
running as.",
      "create": "required",
      "update": "optional"
    },
    "instanceOwnerUid": {
      "type": "integer",
      "description": "The uid of the account the SAP ASE instance is running
as.",
      "create": "readonly",
      "update": "readonly"
    },
    "instanceOwnerGid": {
      "type": "integer",
      "description": "The gid of the account the SAP ASE instance is running
as.",
      "create": "readonly",
```

```

        "update": "readonly"
    },
    "pageSize": {
        "type": "integer",
        "description": "Database page size for the SAP ASE instance."
    },
    "servicePrincipalName": {
        "type": "string",
        "description": "The Kerberos SPN of the database.",
        "create": "optional",
        "update": "optional"
    },
    "dbUser": {
        "type": "string",
        "description": "The username of the database user.",
        "create": "optional",
        "update": "optional",
        "maxLength": 256
    },
    "isqlPath" : {
        "type" : "string",
        "description" : "The path to the isql binary to use for this SAP ASE
instance.",
        "create" : "optional",
        "update" : "optional"
    },
    "credentials": {
        "type": "object",
        "description": "The credentials of the database user.",
        "$ref": "/delphix-credential.json",
        "create": "optional",
        "update": "optional",
        "properties": {
            "type": {
                "type": "string",
                "description": "Object type.",
                "required": true,
                "format": "type",
                "default": "PasswordCredential"
            }
        }
    },
    "discovered": {
        "type": "boolean",
        "description": "True if the SAP ASE instance was automatically
discovered."
    }
}

```

API Cookbook: ASELinkData

This API cookbook recipe describes how to configure your SAP ASE link data using the Delphix Engine API.

```
{
  "name": "ASELinkData",
  "description": "SAP ASE specific parameters for a link request.",
  "extends": {
    "$ref": "/delphix-link-data.json"
  },
  "properties": {
    "config": {
      "type": "string",
      "description": "Reference to the configuration for the source.",
      "format": "objectReference",
      "referenceTo": "/delphix-ase-db-config.json",
      "required": true
    },
    "externalFilePath": {
      "type": "string",
      "description": "External file path.",
      "maxLength": 1024,
      "create": "optional",
    },
    "operations": {
      "description": "User-specified operation hooks for this source.",
      "type": "object",
      "$ref": "/delphix-linked-source-operations.json",
      "create": "optional"
    },
    "mountBase" : {
      "type" : "string",
      "description" : "The base mount point to use for the NFS mounts.",
      "maxLength" : 87,
      "create" : "optional"
    },
    "loadBackupPath": {
      "type": "string",
      "description": "Source database backup location.",
      "maxLength": 1024,
      "required": true
    },
    "loadLocation": {
      "type": ["object", "null"],
      "description": "Backup location to use for loading backups from the
source.",
      "$ref": "/delphix-ase-backup-location.json",
      "create": "optional"
    },
    "dumpCredentials": {
      "type": ["object", "null"],
```

```

    "description": "The credential for the source DB user.",
    "$ref": "/delphix-password-credential.json",
    "create": "optional"
  },
  "sourceHostUser": {
    "type": "string",
    "description": "Information about the host OS user on the source to use
for linking.",
    "format": "objectReference",
    "referenceTo": "/delphix-source-environment-user.json",
    "required": true
  },
  "dbUser": {
    "type": "string",
    "description": "The user name for the source DB user.",
    "create": "optional"
  },
  "dbCredentials": {
    "type": "object",
    "description": "The credentials of the database user.",
    "$ref": "/delphix-credential.json",
    "required": true,
    "properties": {
      "type": {
        "type": "string",
        "description": "Object type.",
        "required": true,
        "format": "type",
        "default": "PasswordCredential"
      }
    }
  },
  "stagingRepository": {
    "type": "string",
    "description": "The SAP ASE instance on the staging environment that we
want to use for validated sync.",
    "format": "objectReference",
    "referenceTo": "/delphix-ase-instance.json",
    "required": true
  },
  "stagingHostUser": {
    "type": "string",
    "description": "Information about the host OS user on the staging
environment to use for linking.",
    "format": "objectReference",
    "referenceTo": "/delphix-source-environment-user.json",
    "required": true
  },
  "stagingPreScript": {
    "type": "string",
    "description": "A user-provided shell script or executable to run prior
to restoring from a backup during validated sync.",
    "maxLength": 1024,

```

```

        "create": "optional"
    },
    "stagingPostScript": {
        "type": "string",
        "description": "A user-provided shell script or executable to run after
restoring from a backup during validated sync.",
        "maxLength": 1024,
        "create": "optional"
    },
    "syncParameters": {
        "type": "object",
        "description": "Sync parameters for the container.",
        "$ref": "/delphix-ase-sync-parameters.json",
        "required": true,
        "properties": {
            "type": {
                "type": "string",
                "description": "Object type.",
                "required": true,
                "format": "type",
                "default": "ASELatestBackupSyncParameters"
            }
        }
    },
    "validatedSyncMode": {
        "type": "string",
        "description": "Specifies the validated sync mode to synchronize the
dSource with the source database.",
        "enum": ["ENABLED", "DISABLED"],
        "default": "ENABLED",
        "create": "optional"
    }
}
}
}

```

API Cookbook: EnvironmentUser

This API cookbook recipe describes how to configure your environment user using the Delphix Engine API.

```
{
  "root": "/resources/json/delphix/environment/user",
  "name": "EnvironmentUser",
  "description": "The representation of an environment user object.",
  "extends": {
    "$ref": "/delphix-user-object.json"
  },
  "nameParent": "environment",
  "properties": {
    "credential": {
      "type": "object",
      "$ref": "/delphix-credential.json",
      "description": "The credential for the environment user.",
      "create": "required",
      "update": "optional",
      "properties": {
        "type": {
          "type": "string",
          "description": "Object type.",
          "required": true,
          "format": "type",
          "default": "PasswordCredential"
        }
      }
    },
    "environment": {
      "type": "string",
      "description": "A reference to the associated environment.",
      "format": "objectReference",
      "referenceTo": "/delphix-source-environment.json",
      "create": "optional"
    },
    "groupId": {
      "type": "integer",
      "description": "Group ID of the user.",
      "create": "optional",
      "update": "optional",
      "minimum": 0,
      "maximum": 4294967295
    },
    "userId": {
      "type": "integer",
      "description": "User ID of the user.",
      "create": "optional",
      "update": "optional",
      "minimum": 0,
      "maximum": 4294967295
    }
  }
}
```

```

    }
  },
  "create": {
    "description": "Create a new EnvironmentUser object.",
    "payload" : {
      "type": "object",
      "$ref": "/delphix-source-environment-user.json"
    },
    "return": {
      "type": "string",
      "format": "objectReference",
      "referenceTo": "/delphix-source-environment-user.json"
    }
  },
  "read": {
    "description": "Retrieve the specified EnvironmentUser object.",
    "return": {
      "type": "object",
      "$ref": "/delphix-source-environment-user.json"
    }
  },
  "update": {
    "description": "Update the specified EnvironmentUser object.",
    "payload": {
      "type": "object",
      "$ref": "/delphix-source-environment-user.json"
    }
  },
  "delete": {
    "payload": {
      "type": "object",
      "$ref": "/delphix-delete-parameters.json",
      "required": false
    },
    "description" : "Delete the specified EnvironmentUser object."
  },
  "list": {
    "description": "Returns the list of all environment users in the system.",
    "parameters": {
      "environment": {
        "type": "string",
        "description": "Limit results to users within the given environment.",
        "format": "objectReference",
        "referenceTo": "/delphix-source-environment.json",
        "mapsTo": "environment"
      }
    }
  },
  "return": {
    "type": "array",
    "items": {
      "type": "object",
      "$ref": "/delphix-source-environment-user.json"
    }
  }
}

```

```
}  
  }  
}
```

API Cookbook: KerberosConfig

This API cookbook recipe describes how to configure Kerberos using the Delphix Engine API.

```
{
  name: "KerberosConfig",
  description: "Kerberos Client Configuration.",
  root: "/resources/json/delphix/service/kerberos",
  singleton: true,
  cliVisibility: [ "DOMAIN", "SYSTEM" ],
  extends: {
    $ref: "/delphix-user-object.json"
  },
  properties: {
    realm: {
      description: "Kerberos Realm name.",
      type: "string",
      create: "required",
      update: "optional"
    },
    kdcs: {
      description: "One of more KDC servers.",
      type: "array",
      create: "required",
      update: "optional",
      minItems: 1,
      items: {
        type: "object",
```

```
        $ref: "/delphix-kerberos-kdc.json"
    }
},
keytab: {
    description: "Kerberos keytab file data in base64 encoding.",
    type: "string",
    format: "password",
    create: "required",
    update: "optional"
},
principal: {
    description: "Kerberos principal name.",
    type: "string",
    create: "required",
    update: "optional"
},
enabled: {
    description: "Indicates whether kerberos has been configured or not.",
    type: "boolean"
}
},
read: {
    description: "Retrieve the specified KerberosConfig object.",

    return: {

        type: "object",
```

```
    $ref: "/delphix-kerberos-config.json"

  }

},

update: {
  description: "Update the specified KerberosConfig object.",

  payload: {

    type: "object",

    $ref: "/delphix-kerberos-config.json"

  }

},

rootOperations: {
  reset: {
    description: "Reset kerberos configuration and disable the feature.",
    payload: {}

  }

}

}
```

API Cookbook: KerberosCredential

This API cookbook recipe describes how to configure Kerberos credentials using the Delphix Engine API.

```
{  
  "name": "KerberosCredential",  
  "description": "Kerberos based security credential.",  
  "extends": {  
    "$ref": "/delphix-credential.json"  
  },  
  "properties": {  
  }  
}
```

API Cookbook: KerberosKDC

This API cookbook recipe describes how to configure KerberosKDC using the Delphix Engine API.

```
{
  name: "KerberosKDC",
  description: "Kerberos Client Configuration.",
  extends: {
    $ref: "/delphix-typed-object.json"
  },
  properties: {
    hostname: {
      description: "KDC Server hostname.",
      type: "string",
      format: "host",
      create: "required",
      update: "optional"
    },
    port: {
      description: "KDC Server port number.",
      type: "integer",
      create: "required",
      update: "optional",
      minimum: 0,
      maximum: 65535,
      default: 88
    }
  }
}
```