

Data Control Tower Home

Data Control Tower

Exported on 02/21/2024

Table of Contents

1	Welcome to the Data Control Tower documentation!	15
2	Quick references	20
3	Release notes	21
3.1	New features	21
3.1.1	Release 14.0.0.....	21
3.1.2	Release 13.0.0.....	21
3.1.3	Release 12.0.0.....	22
3.1.4	Release 11.0.0.....	23
3.1.5	Release 10.0.0.....	23
3.1.6	Release 9.0.0.....	24
3.1.7	Release 8.0.0.....	25
3.1.8	Release 7.0.0.....	25
3.1.9	Release 6.0.0.....	25
3.1.10	Release 5.0.1	26
3.1.10.1	Enhancements	26
3.1.10.2	Custom roles	28
3.1.11	Release 4.0	28
3.1.12	Release 3.0	28
3.1.13	Release 2.2	29
3.1.13.1	Deployment	29
3.1.13.2	APIs.....	29
3.1.13.3	UI.....	29
3.2	Fixed issues.....	30
3.2.1	Release 14.0.0.0 changes	30
3.2.2	Release 13.0.0.0 changes	30
3.2.3	Release 12.0.0.0 changes	30
3.2.4	Release 10.0.1 changes	31

3.2.5	Release 10.0.0 changes	31
3.2.6	Release 9.0.0 changes.....	32
3.2.7	Release 8.0.1 changes.....	32
3.2.8	Release 8.0.0 changes.....	33
3.2.9	Release 7.0.1 changes.....	33
3.2.10	Release 6.0.1 changes.....	34
3.2.11	Release 6.0.0 changes.....	34
3.2.12	Release 5.0.3 changes.....	34
3.2.13	Release 5.0.2 changes.....	34
3.2.14	Release 5.0.1 changes.....	35
3.2.15	Release 3.0.0 changes.....	35
3.3	Supported versions and upgrade matrix	35
4	DCT overview	37
4.1	What is Data Control Tower (DCT)?.....	37
5	Getting started	38
5.1	Planning your deployment.....	38
5.2	Container platform support.....	38
5.2.1	Kubernetes	38
5.2.2	OpenShift.....	38
5.2.3	Docker Compose	38
5.3	Data Control Tower deployment architecture.....	39
5.4	Plan your tagging strategy	39
5.5	Plan your Access Control strategy	40
6	Deployment	43
6.1	Kubernetes	44
6.1.1	Installation and setup for Kubernetes	44
6.1.1.1	Hardware requirements.....	44
6.1.1.2	Kubernetes overview	46
6.1.1.3	Installation requirements (Kubernetes).....	47

6.1.1.4	Installing DCT	47
6.1.2	Ingress setup.....	51
6.1.2.1	Expose proxy HTTP port (80) for non-encrypted traffic.....	51
6.1.2.2	Ingress controller installation and route creation.....	52
6.1.3	Bootstrapping API keys	55
6.1.3.1	Bootstrap first API key.....	56
6.1.4	DCT logs for Kubernetes	58
6.1.5	Admin topics	58
6.1.5.1	Deployment upgrade for Kubernetes.....	59
6.1.5.2	Factory reset DCT for Kubernetes	61
6.2	OpenShift.....	62
6.2.1	Installation and setup for OpenShift.....	62
6.2.1.1	Hardware requirements.....	62
6.2.1.2	Installation requirements (OpenShift)	64
6.2.1.3	Installation process	65
6.2.1.4	Configure Ingress.....	69
6.2.2	OpenShift authentication.....	71
6.2.2.1	Introduction	71
6.2.2.2	Enable OAuth2 authentication	71
6.2.3	DCT logs for OpenShift.....	72
6.2.4	Admin topics for OpenShift.....	72
6.2.4.1	Deployment upgrade for OpenShift	72
6.2.4.2	Factory reset DCT for OpenShift.....	75
6.3	Docker Compose	75
6.3.1	Installation and setup for Docker Compose	76
6.3.1.1	Hardware requirements.....	76
6.3.1.2	Installation requirements (Docker Compose).....	77
6.3.1.3	Unpack and install DCT	78
6.3.1.4	Run DCT.....	78

6.3.1.5	Initial logging configuration	79
6.3.2	Bootstrapping API Keys.....	79
6.3.2.1	Bootstrap first API key.....	80
6.3.3	Custom configuration.....	81
6.3.3.1	Introduction.....	81
6.3.3.2	Bind mounts	81
6.3.4	Docker logs.....	83
6.3.5	Migration topics.....	83
6.3.5.1	Migrate to Kubernetes	84
6.3.5.2	Migrate to OpenShift.....	87
6.3.6	Admin topics for Docker Compose	90
6.3.6.1	Backup DCT on Docker Compose.....	90
6.3.6.2	Deployment upgrade for Docker Compose.....	91
6.3.6.3	Factory reset DCT for Docker Compose.....	93
6.4	Engines: connecting/authenticating.....	94
6.4.1	Introduction.....	94
6.4.2	Truststore for HTTPS	94
6.4.3	Authentication with engine.....	95
6.4.4	HashiCorp vault.....	95
6.4.4.1	Vault authentication and registration	95
6.4.4.2	Token	95
6.4.4.3	AppRole	96
6.4.5	TLS certificates	97
6.4.5.1	Retrieving engine credentials.....	97
6.5	Accounts: connecting/authenticating.....	97
6.5.1	API keys	98
6.5.1.1	Create and manage API Keys.....	98
6.5.2	Username/password	100
6.5.2.1	Password policies.....	102

6.5.2.2	Understanding password policies	102
6.5.2.3	Default password policy	102
6.5.2.4	Changing the password policy	102
6.5.2.5	Disabling local username/password authentication	103
6.5.3	LDAP/Active Directory	103
6.5.3.1	Configuration	103
6.5.3.2	Configure LDAP/Active Directory groups	108
6.5.4	SAML/SSO	111
6.5.4.1	Identity provider setup	111
6.5.4.2	DCT SAML/SSO setup	112
6.5.4.3	Login	113
6.5.4.4	Troubleshooting	113
6.5.4.5	Attributes mapping	114
6.6	Replace HTTPS certificate for DCT	116
6.7	External database support	117
6.7.1	Overview	117
6.7.2	Requirements	117
6.7.3	Setup	118
6.7.3.1	PostgreSQL database setup	118
6.7.3.2	DCT setup	118
6.7.4	Backup and recovery	119
6.7.5	External database migration or upgrade	120
6.7.6	DCT upgrade	120
6.8	DCT data backup, recovery, and migration	121
6.8.1	Example deployment scenarios	121
6.8.2	Prerequisites	122
6.8.3	Directions	122
6.8.3.1	1. Backup	122
6.8.3.2	2. Restore	123

6.8.3.3	3. Additional environment configuration	124
6.9	Exporting DCT logs to Splunk	124
6.9.1	Overview	124
6.9.2	Setting up a Splunk instance.....	124
6.9.3	Enable Splunk log forwarding	124
6.9.4	Search for events in Splunk.....	125
6.10	Generating a support bundle.....	126
6.10.1	Find the “collect_bundle.sh” script	126
6.10.2	Execute the “collect_bundle.sh” script when DCT is running in Kubernetes	127
6.10.3	Execute the “collect_bundle.sh” script when DCT is running in Docker- Compose	127
6.10.4	Find the generated support bundle tar file	128
7	Data governance	129
7.1	DCT administration.....	129
7.1.1	Operations in DCT	129
7.1.1.1	Operations Monitor Bar	132
7.1.2	Tags.....	132
7.1.2.1	Tags management.....	132
7.1.2.2	Administrative tagging	133
7.1.2.3	Tags powering attribute-based Access Control	134
7.1.3	Access Control.....	135
7.1.3.1	Access model overview.....	136
7.1.3.2	Accounts: Manual, LDAP/AD, or SSO/SAML.....	137
7.1.3.3	Access Groups: Creation and account assignment	141
7.1.3.4	Roles: Creation and assignment	147
7.1.3.5	Objects: Refine permission to targeted objects.....	150
7.1.4	VDB templates	157
7.1.4.1	Creating templates	158
7.1.4.2	Importing templates	158

7.1.4.3	Using templates	159
7.1.5	API metering.....	160
7.1.5.1	API metering instructions.....	160
7.1.6	Client telemetry	160
7.2	Central governance workflows	162
7.2.1	Managing engines (Continuous Data).....	162
7.2.1.1	Infrastructure.....	162
7.2.1.2	Engine overview	163
7.2.1.3	Infrastructure connection wizard.....	165
7.2.2	Managing dSources.....	170
7.2.2.1	Managing dSources.....	170
7.2.2.2	dSource overview.....	171
7.2.2.3	Timeflow visibility	171
7.2.2.4	Access auditing.....	172
7.2.2.5	dSource linking.....	172
7.2.3	Managing VDBs.....	173
7.2.3.1	VDB overview	173
7.2.3.2	VDB active timeline.....	173
7.2.3.3	VDB timeline history	174
7.2.3.4	VDB bookmarks	174
7.2.3.5	VDB access	175
7.2.3.6	VDB templates	175
7.2.4	Managing environments (Continuous Data).....	176
7.2.4.1	Global environments list.....	176
7.2.4.2	Manage environments.....	177
7.2.4.3	Edit host details	177
7.2.5	Managing bookmarks.....	178
7.2.5.1	Global Bookmarks List	178
7.2.6	Replication management	179

7.2.6.1	Introduction	179
7.2.6.2	Prerequisites	179
7.2.6.3	User interface	179
7.2.6.4	API.....	181
7.2.6.5	Replication relationship	183
7.3	Insight reports	190
7.3.1	Central governance insights.....	190
7.3.2	VDB Inventory.....	190
7.3.3	dSource Inventory	190
7.3.4	Source Ingestion Metrics	190
7.3.5	Compliance Job Executions report.....	191
7.3.6	Block Storage	191
7.3.7	Activity Audit Log summary	191
7.3.8	Compliance Engine Performance report	191
7.3.9	Data Engine Performance report	191
7.4	Tabular customization support for DCT	192
7.4.1	Introduction	192
7.4.1.1	Key enhancements	193
8	Continuous Data workflows	194
8.1	DevOps TDM	194
8.2	Developer experience	195
8.3	Self-service vs. DCT developer experience	195
8.3.1	Key similarities	196
8.3.2	Key differences	196
8.4	Creating and managing bookmarks	197
8.4.1	Create a Bookmark	197
8.4.1.1	Create a bookmark from an existing snapshot.....	198
8.4.2	Bookmark API Documentation.....	201
8.4.3	Create a Bookmark at the current time for multiple VDBs	201

8.4.4	Creating a bookmark from a chosen timepoint	202
8.5	VDB operations	204
8.5.1	VDB provisioning wizard.....	206
8.5.1.1	Supported database platforms	206
8.5.1.2	Step examples.....	206
8.5.1.3	Limitations.....	212
8.5.2	VDB refresh	212
8.5.2.1	Overview	212
8.5.2.2	User interface.....	212
8.5.3	Active timelines.....	216
8.5.3.1	Active Timeline View	216
8.5.3.2	Additional notes.....	217
8.5.4	Timeline history.....	217
8.5.4.1	Timeline History view	217
8.5.4.2	The Timeline History user interface.....	218
8.5.4.3	Non-active timelines	219
8.5.4.4	Time concepts within the Timeline History Tab	219
8.5.4.5	Timeflow operations.....	220
8.5.4.6	API documentation	222
8.5.5	Locking and unlocking a VDB	223
8.5.5.1	Locking and unlocking a VDB via API.....	225
8.6	Linking Wizard.....	225
8.6.1	Overview	225
8.6.1.1	Limitations in MSSQL dSource linking	226
8.6.1.2	Prerequisites	226
8.6.2	Using the Linking Wizard.....	226
8.6.2.1	1. Select Data Source step	226
8.6.2.2	2. dSource Type step	227
8.6.2.3	3. dSource Configuration step	228

8.6.2.4	4. Data Management step.....	228
8.6.2.5	5. Policies step.....	229
8.6.2.6	6. Hooks step.....	230
8.6.2.7	7. Summary step.....	230
8.6.3	Adding an AppData dSource.....	231
8.6.3.1	Steps effected during the Linking process.....	231
8.6.4	Adding an Oracle Staging Push dSource.....	233
8.6.4.1	Prerequisites.....	233
8.6.4.2	Linking process.....	234
8.6.4.3	Completion.....	237
8.6.5	Adding a MSSQL Staging Push dSource.....	237
8.6.5.1	Prerequisites.....	237
8.6.5.2	Linking process.....	237
8.6.5.3	Completion.....	239
9	Continuous Compliance workflows.....	240
9.1	Listing and searching compliance jobs.....	240
9.2	Consolidated operations (intelligent syncing).....	241
9.3	Managing engines (Continuous Compliance).....	241
9.3.1	Engine overview.....	242
9.3.2	Engine-based operations access.....	243
9.4	Compliance jobs.....	243
9.4.1	Job UI.....	244
9.4.1.1	Global compliance jobs list.....	244
9.4.1.2	Compliance job overview.....	244
9.4.1.3	Access tab.....	245
9.4.1.4	Execution history tab.....	245
9.4.1.5	Execution details.....	246
9.4.2	Copy job.....	247
9.4.2.1	User interface documentation.....	248

9.4.2.2	API documentation	248
9.4.3	Execute job	250
9.4.3.1	User interface	252
9.4.4	Migrate job	252
9.4.5	Delete job.....	253
9.5	Read-only algorithms.....	254
9.5.1	Overview	254
9.5.1.1	Feature Limitations.....	254
9.5.1.2	Algorithms page.....	254
9.5.1.3	Algorithm Details Overview page.....	255
9.5.1.4	Algorithm Revisions page.....	256
10	Hyperscale Orchestrator UI	258
10.1	Overview	258
10.2	Implementation introduction.....	259
10.2.1	Truststore for HTTPS	259
10.2.2	Authentication with Hyperscale Orchestrators.....	259
10.2.3	Hyperscale deployment type.....	260
10.2.4	Editing and unregistering Hyperscale Orchestrators.....	260
10.3	Managing Hyperscale objects.....	261
10.3.1	Introduction	261
10.3.2	Hyperscale executions	262
10.3.3	Hyperscale jobs.....	262
10.3.4	Hyperscale job engine selection	264
10.3.5	Hyperscale job table configuration.....	265
10.3.6	Executing Hyperscale jobs	266
10.3.7	Creating Hyperscale jobs	268
10.3.8	Hyperscale Compliance Engines	271
10.3.9	Hyperscale mount points	273
10.3.10	Hyperscale connector configurations	275

11	Integrations	277
12	DCT concepts.....	278
12.1	Introduction	278
12.2	Concepts	278
12.2.1	Virtual Database (VDB) groups	278
12.2.2	Comparing Self-Service containers to VDB groups	279
12.2.3	Bookmarks	280
12.2.4	Jobs	280
12.2.5	Tags	281
12.2.6	Tag-based filtering.....	281
12.3	Nuances.....	282
12.3.1	Stateful APIs.....	282
12.3.2	Local data availability	282
12.3.3	Engine-to-DCT API mapping.....	282
12.3.4	Local references to global UUIDs	282
12.3.5	Environment representations.....	282
12.3.6	Supported data sources/configurations	283
12.3.7	Process feedback	283
13	DCT Toolkit.....	284
13.1	Introduction	284
13.2	Compatibility	284
13.3	New features	284
13.3.1	Version 1.2.0	284
13.4	Installation and setup	285
13.4.1	Installation.....	285
13.4.2	Setup.....	286
13.5	Usage guide.....	288
13.5.1	Examples	288
13.5.1.1	--json/-js and --csv/-cs option examples	290

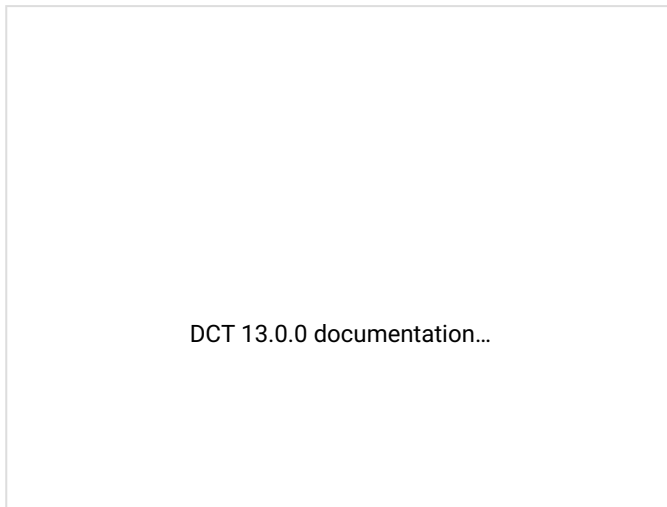
13.6	API key encryption	291
13.6.1	Overview	291
13.6.2	Implementation	291
13.6.3	Backward compatibility	292
13.6.4	Example	292
13.7	Configure multiple DCT instances in dct-toolkit	292
13.7.1	Overview	292
13.7.2	Example	292
13.8	Logging	293
14	Developer resources	294
14.1	API requests and reporting	294
14.1.1	Introduction	294
14.1.2	Engines	294
14.2	API references	295

1 Welcome to the Data Control Tower documentation!

This information explains how to deploy Data Control Tower (DCT), use its features, or tune its configurations for optimal performance. The content has been organized into several categories, available from the lefthand navigation.

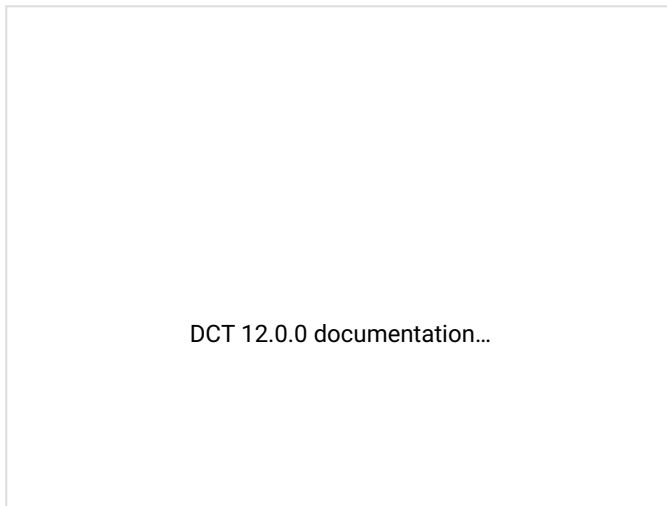
List of Data Control Tower documentation versions in PDF format.

-



[\(see page 15\)](#)

-





[\(see page 15\)](#)

-

- [DCT 11.0.0 documentation...](#) [\(see page 15\)](#)
- [DCT 10.0.0 documentation...](#) [\(see page 15\)](#)
- [DCT 9.0.0 documentation....](#) [\(see page 15\)](#)

- [DCT 8.0.0 documentation....](#) [\(see page 15\)](#)
- [DCT 7.0.0 documentation....](#) [\(see page 15\)](#)
- [DCT 6.0.0 documentation....](#) [\(see page 15\)](#)

- [DCT 5.0.0 documentation....](#) [\(see page 15\)](#)
- [DCT 4.0.0 documentation....](#) [\(see page 15\)](#)
- [DCT 3.0.0 documentation....](#) [\(see page 15\)](#)

-  DCT 2.2.0 documentation... [\(see page 15\)](#)
-  DCT 2.0.0 documentation... [\(see page 15\)](#)

2 Quick references

- [Overview](#)¹
- [Deployment](#)²
- [New features](#)³
- [Fixed issues](#)⁴
- [Concepts](#)⁵

1 <https://dct.delphix.com/docs/latest/dct-overview>

2 <https://dct.delphix.com/docs/latest/deployment>

3 <https://dct.delphix.com/docs/latest/new-features>

4 <https://dct.delphix.com/docs/latest/fixes>

5 <https://dct.delphix.com/docs/latest/dct-concepts>

3 Release notes

This section is used to learn what the newest version of Data Control Tower has to offer. In addition, the fixed and known issues per version are detailed.

- [New features](#) (see page 21)
- [Fixed issues](#) (see page 30)
- [Supported versions and upgrade matrix](#) (see page 35)

3.1 New features



Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁶ or [OpenShift](#)⁷.

3.1.1 Release 14.0.0

- **New APIs for PostgreSQL**
New APIs have been added, providing the function to create, update, or delete database APIs for PostgreSQL.
- **Bug fixes and improvements**
This release focuses on bug fixes and general improvements.

3.1.2 Release 13.0.0

- **Algorithm centralization**
To enable better visibility into compliance, masking algorithms across all connected Continuous Compliance Engines can now be centrally viewed in DCT.
- **Expanded source linking**
Over the past few years, Staging Push has been introduced, allowing users to bring their own data to

⁶ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁷ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

Continuous Data. The DCT UI now supports linking for Oracle, SQL Server using Staging Push. In addition, ASE and SDK connectors such as Postgres, Db2, MySQL, HANA, and MongoDB.

- **Hyperscale Compliance UI enhancements**
Several improvements have been added, including enabling the management of Continuous Compliance Engine assignments of Hyperscale Jobs.
- **DCT Toolkit enhancements**
Based on user feedback, several improvements have been made to the fully supported CLI for DCT. Notably, `dct-toolkit` allows users to provide engine names as values for 'id' options.
- **Continuous Data jobs progress**
Users can track job progress in real-time through the Operations UI.
- **Improved Replication visibility**
Replication, used for data movement, disaster recovery protection, Sensitive Data Distribution, and more, requires setting up relationships between Continuous Data Engines. To provide better visibility on these relationships, they have been added to DCT. To further simplify this user experience, the target Continuous Data Engine and dataset are now linked together.
- **Engine Performance trend lines**
Historical trend lines have been added to the recently introduced Engine Performance reports.

3.1.3 Release 12.0.0

- **Delphix Kubernetes driver**
Virtual databases (vDB) can be provisioned into containers through Helm charts and `kubectl` commands using a new, pre-built DCT integration. Teardown of these containers can also be automated, enabling ephemeral infrastructure for Kubernetes-based applications.
- **Engine performance reports**
New Continuous Data and Compliance reports have been added to improve the central management of engine infrastructure performance. You can now view allocated resources and track performance for all connected engines, including disk latency, network latency, and throughput.
- **Hyperscale Compliance jobs**
DCT now has support for creating, managing, and executing Hyperscale Compliance jobs. The concept of Engine Pools have been introduced along with several other enhancements to refine the overall experience.
- **User experience improvements**
Several enhancements including customizing the columns across all tables, increased color contrast for better accessibility, simplified navigation across tabs, and a variety of visual style improvements.
- **Replication tag inheritance**
DCT now offers API-based simplified tag management for parent/child replication deployments, optionally syncing source tags to a target replica.
- **Oracle and SQL Server source linking**
Linking sources for Oracle and SQL Server is now supported in the user experience, along with creating bookmarks.
- **Terraform provider**
A new Terraform resource facilitates creating and deleting Oracle dSources. This expands on the existing data source support with the provider.

3.1.4 Release 11.0.0

- **dSource linking support (API only)**

The ability to create dSources for MSSQL (single instance and cluster), Oracle Staging Push, and MSSQL Staging Push has been added. In addition, dSource linking for older engine versions back to 6.0.7 is now supported.

- **Last Refresh column added to the VDB tab**

A new column called Last Refresh has been added to the VDB tab, which shows the last refresh date/time of a VDB.

- **download.delphix.com⁸ enhancements for container-based upgrades**

Changes have been made to download.delphix.com⁹, making it easier to build installation and upgrade automation for DCT using repository-management products like JFrog Artifactory.

- If you are installing DCT for the first time, please follow the standard installation documentation.
- If you have an existing installation of DCT you must take the following steps to perform upgrades:
 - Update the registry URL to the new address (<http://dct.download.delphix.com> and/or <http://hyperscale.download.delphix.com>) in the values.yaml file.
 - Update the credentials with the new permanent credentials in the values.yaml file.
 - If your organization leverages an allow (white) list, please update the list with the new repository URL(s).

3.1.5 Release 10.0.0

- **dSource linking support (API only)**

We have added the ability to create dSources for Oracle, ASE, and Postgres via API. This is currently available for Delphix Continuous Data Engines on version 8.0 and above. Future releases will include full support for all data sources and inclusion of older engine versions.

- **Hyperscale Compliance UI (read-only)**

DCT has released a Hyperscale UI under controlled availability. Please reach out to your account team to activate this feature. This enhancement includes the ability to:

- Register Hyperscale Orchestrators with DCT.
- Visualize Hyperscale infrastructure configuration parameters such as clusters of engines and mount points.
- Report on Hyperscale Jobs and Job executions, including near-realtime process feedback.

- **Compliance Job Executions report**

For connected Continuous Compliance Engines, DCT now has a new report under DCT Insights that

⁸ <http://download.delphix.com/>

⁹ <http://download.delphix.com/>

displays all recent job executions (including DCT initiated, engine initiated, and Hyperscale initiated jobs) as well as relevant compliance metrics.

- **Provisioning wizard enhancements**

The DCT provisioning wizard now supports provisioning Oracle multi-instance (multi-tenant) and ASE VDBs.

- **Operations page**

The DCT Operation page now includes initiator details, which provides visibility of what user/account initiated a job.

- **Support for creating a bookmark at point-in-time**

DCT now has the ability to create bookmarks on VDBs under a specified point-in-time via both API and UI. This feature is only compatible for databases that support LogSync and have it enabled.

- **VDB lock/unlock**

Users now have the ability to lock a VDB, preventing any sort VDB-based activity (refresh, start, stop, delete, and create bookmark).

- **Custom hook support in the DCT UI**

Delphix users now have the ability to add custom hooks to a VDB at the time of provision (as part of the provision wizard UI) and update them under a VDB details page, so that they will execute during refresh operations.

3.1.6 Release 9.0.0

- **Jenkins support**

An official Jenkins plugin is now available for Data Control Tower, joining the existing ServiceNow and Terraform integrations. This plugin helps automate the use of data in your CI/CD pipelines and includes support for provisioning and destroying VDBs.

- **DCT Toolkit**

We've launched a new command line interface (CLI) offering, the DCT Toolkit. The DCT Toolkit allows for remote operation and control of Data Control Tower via your local terminal. Over time, this will replace the previously available DxToolkit.

- **Infrastructure wizard**

The create infrastructure flow now supports adding Windows standalone and cluster hosts, as well as Linux cluster hosts. This is expanded from the previous flow that supported standalone Linux hosts. This differs from the workflow on the Continuous Data engine in that it represents a single place to create environment host connections for all connected Delphix engines.

- **Provisioning wizard**

The data provisioning workflow now includes additional database types, including Oracle Single Instance Multi-tenant, Oracle Multiple Instance Single Tenant, and SQL Server Multiple Instance Single Tenant.

- **Operations dashboard improvements**

Previously, all users could view all operations run by any user on the Operations page. Role-based access can now be provided to different users, showing only operations the user has access to. Simplified text is now used to define the Type column, as well as the Type and Engine Name in the details page.

- **Activity Audit Log Summary report**
The Activity Audit Log Summary provides a high-level audit log summary capturing the utilization of DCT by displaying user activity and the historical count of actions executed within the platform.
- **Replication mappings**
This feature helps users differentiate between replicated objects and original objects, in case of master and replicated engine, both are registered with DCT.

3.1.7 Release 8.0.0

- **Operations dashboard**
Monitor and manage enterprise data activities in real time using a new central view. This provides visibility to the current status across the full complement of Delphix transactions, including provision, refresh, teardown, and compliance jobs.
- **Provisioning wizard enhancements**
The data provisioning wizard has been expanded to support additional types, including Oracle Single Instance Linked CDBs and Microsoft SQL Server Single Instance workflows. This will now allow you to provision more data types directly from Data Control Tower.
- **Advanced search tags support**
The advanced search capabilities now support all user-generated tags. You can use personalized tags related to your unique business needs to refine your search results, such as team names or other specific data points.

3.1.8 Release 7.0.0

- **Provision VDB UI**
Extending the Developer Experience capability in DCT, users can now provision single-tenant Oracle databases from the user interface using an intuitive wizard workflow.
- **Refresh VDB UI enhancements**
The VDB list can now be opened in a searchable, paginated list selector from within a dialog by clicking the select button in the input. Additionally, refreshing a VDB by a bookmark is now available.
- **VDB template import**
Importing and removing imported VDB templates from connected engines is now an available action from the "VDB Config Templates" page.
- **Environment details enhancements**
Managing cluster environment infrastructure has been made easier with the ability to edit host details directly from the page.

3.1.9 Release 6.0.0

- **Developer Self-Service UI**
Developers and admins now have the ability to centrally orchestrate common Continuous Data and developer operations from the DCT UI. This includes the ability to refresh, rewind, bookmark, and bookmark share (refresh to relative). This functionality also exposes the notion of **time flows** (non-

active timelines), which is a critical tool for viewing past work on a VDB, such as the chronology of test results.

- **Central compliance orchestration**
The compliance job UI now enables job orchestration and reporting. This includes Job Copy and Execute functions as well as a complete historical job execution log within each compliance job's details view.
- **Bookmark UI**
Developers and admins now have added visibility of bookmarks, both globally and contextualized, to individual VDBs. These visualizations are dual purpose; for administrators, these screens help with reporting and tagging on bookmarks, while for developers, these screens act as a catalog of actionable data references.
- **Global Bookmark List**
View all bookmarks across your entire connected Delphix ecosystem. This screen will show bookmarks for both single VDBs and VDB groups.
- **VDB Bookmark List**
See all bookmarks tied to this individual VDB. This is helpful for sharing bookmarks with team members who have a compatible VDB (same parent and provision point).
- **Environment details page**
Users can now orchestrate common environment actions via the DCT UI including enable, disable, environment refresh, and delete, as well as editing host details. Note, editing host details is only applicable to standalone environments at this time.
- **Access visibility**
Object detail pages will include an access tab that provides visibility to user access and the associated permissions for each user. This is a critical enabler for permissions visibility and auditing.
- **Copy/delete functionality on role scopes**
Scoped roles can now be copied and deleted within the DCT UI. This will enable easier administration, especially around the use of custom roles, as admins can now copy and modify new roles from templates.
- **External Postgres DB support**
DCT now supports the use of an external Postgres database to house DCT metadata. Previously, DCT supplied and managed its own database, requiring persistent storage within the container platform.

3.1.10 Release 5.0.1

3.1.10.1 Enhancements

- **Data scoped Access Group**
 - **Enhancement in roles**
Associated permissions in roles are changed from 'string' type to 'permission object' type. For details, see the Role schema in the [API References \(see page 295\)](#).
 - **Custom roles**
In addition to the 5 pre-seeded fixed roles (Admin, Monitoring, DevOps, Masking, and Owner), DCT provides flexibility to create new custom roles as per user need. Users (Accounts) can create new custom roles by encapsulating any combination of permissions. The custom roles

can be configured through a UI configuration screen (screenshot below), in addition to a set of APIs to manage roles. For details, see the [API References \(see page 295\)](#).

- **Updates to existing RBAC model**

For better usability and allow to set more granular permissions there are following enhancements in the RBAC model:

- **Renamed Access Group "Policy" to Access Group "Scope"**
- **Renamed the following APIs related to Access Group actions**

- **Add scope** to an Access Group

POST: /access-groups/{accessGroupId}/policies → POST /access-groups/{accessGroupId}/scopes

- **Remove scope** from Access Group

DELETE /access-groups/{accessGroupId}/policies/{policyId} →
DELETE /access-groups/{accessGroupId}/scopes/{scopeId}

- **Get** Access Group scope

GET /access-groups/{accessGroupId}/policies/{policyId} → GET /access-groups/{accessGroupId}/scopes/{scopeId}

- **Update** Access Group scope

PATCH /access-groups/{accessGroupId}/policies/{policyId} →
PATCH /access-groups/{accessGroupId}/scopes/{scopeId}

- **Add object tags** to Access Group scope

POST /access-groups/{accessGroupId}/policies/{policyId}/object-tags → POST /access-groups/{accessGroupId}/scopes/{scopeId}/object-tags

- **Remove object tags** from Access Group scope

POST /access-groups/{accessGroupId}/policies/{policyId}/object-tags/delete → POST /access-groups/{accessGroupId}/scopes/{scopeId}/object-tags/delete

- **Add objects** to Access Group scope

POST /access-groups/{accessGroupId}/policies/{policyId}/objects → POST /access-groups/{accessGroupId}/scopes/{scopeId}/objects

- **Remove objects** from Access Groups scope

POST /access-groups/{accessGroupId}/policies/{policyId}/objects/delete → POST /access-groups/{accessGroupId}/scopes/{scopeId}/objects/delete

- **Renamed the "everything" flag to "scope_type"**

In order to make it more understandable, we have renamed the everything flag to scope_type. There are three possible values for scope_type i.e. SIMPLE, SCOPED and ADVANCED. The

value SIMPLE corresponds to everything=true and SCOPED corresponds to everything=false. The value ADVANCED for scope_type is new enhancement to setting permissions which allows users to set permissions (e.g. READ, DELETE) for an object. There is more information about ADVANCED scope in next section.

- **Access Group Scope: Advanced scope type**
In Add objects to access group scope API, now user can define permissions level checks as well for an object. For example, earlier when object_id and object_type are provided in request payload, all permissions that are defined in scope are applied to this object. But now user can define specific permissions.
- **Masking jobs**
 - CRUD APIs, COPY, Connectors CRUD
- **Masking job execution**
 - Connector Credentials
 - Execution API

3.1.10.2 Custom roles

- Accounts can create new instances of role encapsulating any combination of permission.
- Role name must be unique.
- Custom roles can be updated. Accounts can add or remove permissions to/from the custom roles.
- Custom roles can be deleted. (If they are not associated with any Access Group).

3.1.11 Release 4.0

- Environment Overview List
- Un-virtualized Source Sizing Report
- Global VDB Templates
- Scoped Access Control
- LDAP/AD and SAML/SSO Configuration UI

3.1.12 Release 3.0

- Cluster Node (RAC) management APIs
- Ability to disable username/password authentication globally
- LDAP/Active Directory groups
- CDBs/vCDBs APIs
- VDB Provisioning / update for EDSI (AppData) platforms
- Engine registration wizard
- Access Groups Management UI
- Compliance Engine Management

3.1.13 Release 2.2

3.1.13.1 Deployment

- Introducing Kubernetes and OpenShift support

3.1.13.2 APIs

- Registration of Continuous Compliance Engines
- Masking Connectors
- “Move Masking Job”
- Masking of mainframe objects
- Provisioning enhancements for Oracle multi-tenant and RAC
- LDAP/Active Directory authentication
- Password management
- Initial access management by Permissions, Roles, Policies, and Access Groups (permissions applied to all objects of a type e.g. Stop VDB permission on all VDBs)
- Distributed tracing and logging (Trace ID propagated down call stack)
- Bulk delete of tags

3.1.13.3 UI

- Continuous Data
 - Added tag support to the Infrastructure page
 - New dSources page
 - New VDBs page
- Insights
 - Added an export behavior to the Storage Summary report
 - New dSource Inventory report
 - New VDB Inventory report
- Admin
 - New Accounts page

3.2 Fixed issues

3.2.1 Release 14.0.0.0 changes

Bug Number	Description
APIGW-7090	Fixed an issue where the username was not displayed in the top navigation bar, if logging in using anything other than username/password.

3.2.2 Release 13.0.0.0 changes

Bug Number	Description
APIGW-4971	Fixed a problem when <code>VDB_GROUP_REFRESH</code> job's update time did not get updated on receiving job updates for the underlying <code>VDB_REFRESH</code> jobs.
APIGW-6705	Fixed an installation problem when masking database migration runs out of memory for large tables.

3.2.3 Release 12.0.0.0 changes

Bug Number	Description
APIGW-6240 APIGW-6241	Fixed a memory exhaustion issue that caused frequent HTTP 500 errors.
APIGW-5052	Fixed an issue when viewing the Summary screen after provisioning a vDB in DCT where the dSource summary name was showing the vDB name.
APIGW-5511	Added the ability to unregister engines in the UI.
APIGW-5585	Fixed an issue where users were unable to logout following SSO token expiration.

3.2.4 Release 10.0.1 changes

Bug Number	Description
APIGW-5406	Fixed an issue where the VDBGGroup update API gives an internal error if vdb_ids are the same as ones already present in VDBGGroup.
APIGW-5419	Fixed an issue where the VDBGGroup delete API gives a 404 error if underlying VDBs are already deleted or not present.
APIGW-5418, APIGW-5517	Fixed an issue where VDB provisioning via DCT GUI fails with, "There was an error trying to process your request."
APIGW-5570	Fixed an issue preventing the link of a Postgres database as a dSource.
APIGW-5571	Fixed an issue where the Provision VDB -> dSource Listing would not show more than 25 items.
APIGW-5574	Hyperscale datasets tables-or-files list and search APIs now return all instead of filtering by dataset.
APIGW-5578	Fixed an issue where Pagination was not working for some of the pages in a testing environment.

3.2.5 Release 10.0.0 changes

Bug Number	Description
APIGW-3931	Fixed an error found when loading the LDAP config page in the UI, if the LDAP config domains have empty values.
APIGW-3961	Fixed the issue where the Environment link to the dSource detail view is broken.
APIGW-4270	Cleaned up pending jobs left by deleted engines.

Bug Number	Description
APIGW-5056	Addressed VDB provisioning failures where there were a lot of concurrent requests.

3.2.6 Release 9.0.0 changes

Bug Number	Description
APIGW-3772	Replicated VDBs/dSources are now identified in the Provisioning wizard.
APIGW-3931	Fixed a null pointer exception during LDAP configuration without a domain.
APIGW-3979	Fixed an issue where the Import VDB Configuration templates dialog was showing the engine ID in error messages. It is now changed to show engine names.
APIGW-3983	Fixed an issue where a new masking job could not be started from DCT when the previous job was cancelled on the masking engine.
APIGW-4009	Fixed an issue where the first and last name will be cleared if incorrect names were entered for SSO.
APIGW-4010	Fixed an issue in the UI where first and last name attributes cannot be reset.

3.2.7 Release 8.0.1 changes

Bug Number	Description
APIGW-4324	Fixed an issue where users who upgraded to DCT 8.0.0 were not able to interact with the UI or connect to the GraphQL service container.
APIGW-4317	Fixed an issue where an error would occur when searching for a VDB in the relative refresh UI.

3.2.8 Release 8.0.0 changes

Bug Number	Description
APIGW-3764	Removed THE requirement on setting credentials if a masking job execution happens on the origin engine.
APIGW-3771	Allows the policy name to be empty when provisioning a VDB.
APIGW-3783	Allows for an existing ImagePullSecret to be provided to to pull docker images.
APIGW-3985	Fixed the "VDB Container is part of a container" error while refreshing from bookmark directly on the VDB > Bookmark tab.
APIGW-3990	Fixed the broken view for a bookmark that has multiple VDBs on the Data > Bookmark tab.

3.2.9 Release 7.0.1 changes

Bug Number	Description
APIGW-3592, APIGW-3594	Previously, a non-admin user that was granted access to a VDB, but not its environment, would get an error accessing the VDB overview. A fix has been implemented to show that the access error is with the environment and not the VDB.
APIGW-3775	Fixed an issue where refreshing from the bookmark wizard was not showing compatible bookmarks.
APIGW-3831	Fixed a certificates import failure if the truststore is on OpenShift.

3.2.10 Release 6.0.1 changes

Bug Number	Description
APIGW-3460	Fixed a request timeout issue.
APIGW-3395	Fixed an issue where the refresh wizard did not update snapshots when selecting different datasets.

3.2.11 Release 6.0.0 changes

Bug Number	Description
APIGW-3223	Fixed an issue where DCT failed to get info from detached dSources.

3.2.12 Release 5.0.3 changes

Bug Number	Description
APIGW-3344	Fixed an issue causing provision failure from RAC dSource to non-RAC target.

3.2.13 Release 5.0.2 changes

Bug Number	Description
APIGW-2979	VDB refresh will no longer fail if the refresh target name is not unique.
APIGW-2981	Fixed an issue where all the Compliance jobs and source jobs on the engine will be deleted when a Compliance engine is unregistered.

3.2.14 Release 5.0.1 changes

Bug Number	Description
APIGW-2463	The default docker-compose.yaml file is now provided with log size and rotation configured for all containers.
APIGW-2735	Fixed an issue where DCT migration failed with "could not create unique index environments_host_pkey".
APIGW-2828	Helm chart now allows cronjob resource limits to be set via the values.yaml.

3.2.15 Release 3.0.0 changes

Bug Number	Description
APIGW-1785	Fixed an issue where Nginx sometimes failed to start after a server restart.

3.3 Supported versions and upgrade matrix

Data Control Tower has minimum engine versions that are actively tested against to ensure optimal interoperability. Please ensure that all connected engines meet the version requirements:

Delphix Engine	Version
Continuous Data	6.0.0.1 or higher
Continuous Compliance	6.0.13.0 or higher

Users can upgrade directly between DCT versions without needing an interim step (i.e., upgrading to a median version before upgrading to the latest).

Version	Release date	Can upgrade to
2.0.0	Jun 27th, 2022	2.1.0 - 13.0.x


Version	Release date	Can upgrade to
2.1.0	Sep 8th, 2022	2.2.0 - 13.0.x
2.2.0	Oct 17th, 2022	3.0.0 - 13.0.x
3.0.0	Dec 15th, 2022	4.0.0 - 13.0.x
4.0.0	Jan 19th, 2023	5.0.0 - 13.0.x
5.0.x	Feb 16th, 2023	6.0.0 - 13.0.x
6.0.x	Mar 29th, 2023	7.0.0 - 13.0.x
7.0.x	May 9th, 2023	8.0.0 - 13.0.x
8.0.x	Jun 22nd, 2023	9.0.x - 13.0.x
9.0.x	Aug 1st, 2023	10.0.x - 13.0.x
10.0.x	Sep 13th, 2023	11.0.x - 13.0.x
11.0.x	Oct 23rd, 2023	12.0.x - 13.0.x
12.0.x	Nov 30th, 2023	13.0.x
13.0.x	Jan 24th, 2024	N/A

Where x represents patch version releases (i.e. 5.0.1, 5.0.2, etc.).

4 DCT overview

4.1 What is Data Control Tower (DCT)?

Today's application and data landscape is an increasingly complex ecosystem of hosting architectures, often represented by a multi-cloud landscape coupled with an explosion of different platforms and services. This fragmented picture of heterogeneous silos makes data governance, automation, and compliance a herculean, if not, an impossible task.



Data Control Tower (DCT) is an enabling Delphix platform that introduces a data mesh to unify data governance, automation, and compliance across all applications and cloud platforms.

Data governance is achieved through operational control and visibility of test data across multicloud applications, databases, environments, and releases. DCT brings data cataloging, tagging, and data access controls for central governance of all enterprise data, while providing the right data at the right time to development teams.

Data automation at CI/CD speed and enterprise scale is easier and more powerful, by combining **DCT with Continuous Data**. A unified API gateway, self-service automation tools, and plug-and-play DevOps integrations streamline the initial configuration and day-to-day workflows.

DCT with Continuous Compliance provides robust data compliance in lower environments, all while reducing costs and enabling fast, quality software development.

5 Getting started

5.1 Planning your deployment

Data Control Tower (DCT) represents a Delphix-wide control plane. It simultaneously powers data governance, automation, and compliance workflows to enable the efficient operation of a broad, complex Delphix deployment at scale. In order to deliver scalability, service-level performance tuning, and robust resiliency, DCT leverages container technology to deliver a bespoke experience for administrative teams based on their own internal guidelines.



Before starting a DCT deployment, please contact your enterprise IT organization to determine what container platforms, configurations, and policies apply for container-hosted applications. **It is helpful to include a container administrator as part of the DCT install process.**

5.2 Container platform support

Data Control Tower (DCT) supports the most popular distributions of Kubernetes and OpenShift. If you do not see your distribution or platform of choice, please reach out to your account team for more details.

5.2.1 Kubernetes

DCT currently supports all popular deployment models of Kubernetes as long as the service runs a minimum of Kubernetes 1.25 and above. This includes Amazon Elastic Kubernetes Service (Amazon EKS), Azure Kubernetes Service (AKS), and beyond.

5.2.2 OpenShift

DCT also supports all popular deployment models of OpenShift as long as the service runs a minimum version of 4.12 or above. This includes Red Hat OpenShift on IBM Cloud and any other cloud provider's service.

5.2.3 Docker Compose

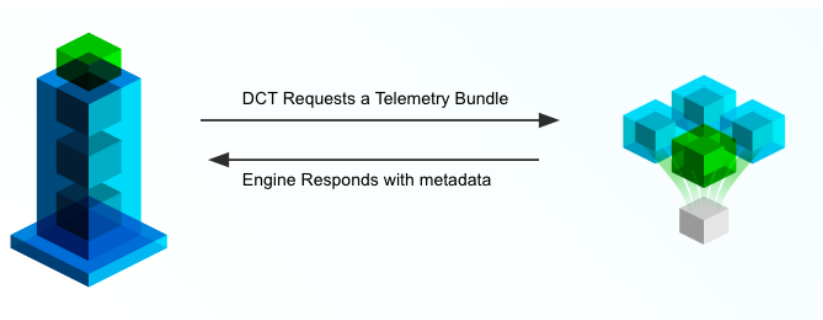
DCT supports Docker Compose but only recommends using this platform for testing/non-production purposes due to the inherent limitations to deployment scalability. Note - DCT has documentation on migrating deployments from Docker Compose to Kubernetes and OpenShift.

5.3 Data Control Tower deployment architecture

Whether an organization wants to deploy a Data Control Tower (DCT) per business unit (organizational silos), per network (datacenter-specific DCT), or globally (**the most common option**), DCT can adapt to many deployment scenarios.

- Delphix recommends to deploy a single, global DCT for all Delphix Engines, for the purposes of achieving a single control plane and data governance solution.

DCT-based communication is lightweight, requiring simple commands or a small telemetry payload to facilitate most workflows. The below graphic demonstrates this style of communication:



DCT simply logs into the engines as a user would and leverages engine APIs to perform commands or extract metadata.

- DCT requires HTTP/HTTPS to facilitate communication with engines and requires ports 80/443 to reach engines in other networks.

- DCT does not directly interface with business-critical databases, it will only communicate with engines to perform operations and inquire about system statuses. The Delphix Engine, which is generally co-located with your data, does all the heavy lifting.

5.4 Plan your tagging strategy

DCT tags serve as the Delphix-wide business metadata system. These **Key:Value** pairs can be applied to any object and used for search and filter in virtually every DCT workflow, from automation to administration, all the way to access control.

- It is paramount to develop a tagging strategy prior to deployment in order to develop a scalable metadata solution.

Some examples of popular tagging strategies:

Theme	Sub-topics	Tag (Key:Value) example
Owner	Application, Business, Project, Team (scrum, QA,...)	(Owner: Finance App), (Owner: AppTeam Alpha), (Owner: John Doe), ...
Application		(Application: Alpha)
Environment		(Environment: Non Production)
Location	Data Center, Region, Name, Cloud	(Geo: West Coast), (Data Center: Azure WC), ...

In addition to designing which tags to implement, please consider who will have access to creating tags (i.e. developer vs admin-only, etc.), which will impact how teams are able to collaborate with one another.

Also, Delphix recommends that the DCT administrative team creates Delphix-wide documentation on these tagging standards to reduce the risk of deviation.


5.5 Plan your Access Control strategy

DCT implements a model found in other types of software called **Attribute Based Access Control (ABAC)**. This model is incredibly flexible, but requires detailed configurations to perfect your use cases. In DCT's model, there are four entity types (defined below). Familiarize with each entity, as they are the foundational blocks of DCT's ABAC model.

Entity	Description	Managed by
Accounts (aka Users)	A single or shared user who can authenticate with DCT (UI or API).	Create manually or via Identity Provider (IdP), such as SSO or LDAP. Accounts are independent of Delphix Engines.

Access Groups	A collection of accounts that share one or more characteristics, such as a Team or Permission set. Equivalent to an Active Directory group.	Manually created. Populated manually or via the <code>login_groups</code> tag.
Roles and Permissions	The collection of read, write, and delete permissions forms a reusable, named role.	Some roles are provided out of the box, but Admins can build their own from the available permissions. Individual permissions are immutable.
Objects	Units, such as VDBs, Bookmarks, and Environments, that are managed across the Delphix Platform.	Automatically identified by DCT from the connected engines. Assigned to Roles via various models. The CD and CC Engines supply these objects.

Each entity is linked to another through manual or automated assignment. A manual (or direct) assignment is a good approach for early implementations, however, that can be challenging to maintain as teams grow. As an alternative, tagging is recommended to perform automatic assignments based on your custom configuration. The below diagram shows how each entity is linked together. The directions below start with **Accounts creation to Access Groups** with **Role assignments**, and finish with **Object mappings**.



Understanding your team structure is imperative to identify the best access model. Usually, organizations have existing groupings defined in their Identify Provider (IdP). These groups are typically organized in one of two ways:

- A team dedicated towards a central goal (such as a Product Development team).
- A group of individuals with similar permissions (such as Security Administrators).


Understanding the purpose of each group should be a guide in how the Roles and Permissions are designed. For example, the Alpha product development team might have full permission to manage existing VDBs and create new bookmarks for their team’s “Alpha” objects. On the other hand, Security Admins might have sweeping read and disable access across the entire platform to ensure compliancy. Iterating through each Access Group and designing custom, but re-useable roles, based on the [Principle of Least Privilege](#)¹⁰, will produce a streamlined rollout.

¹⁰ https://en.wikipedia.org/wiki/Principle_of_least_privilege


6 Deployment

The articles in this section will explain all of the required steps to deploy DCT on your container platform of choice.

Data Control Tower is a container-based architecture that is currently certified with **Kubernetes** and **OpenShift** to align with common enterprise container standards. The DCT architecture is comprised of multiple micro-services that are each run on individual pods. This lends DCT to be a highly flexible deployment by enabling customers and IT organizations to enact their own backup, scaling, and resiliency standards associated with hosting container-based applications. Below is an architectural diagram of all the services that make up DCT, as well as the persistent storage for maintaining relationship metadata.



DCT is **multi-cloud enabled**, which means that a single DCT instance can be deployed to orchestrate (via HTTPS) Continuous Data and Continuous Compliance workloads with Delphix Engines located in other networks. Alternatively, DCT can be localized to engines located within a network. DCT is a lightweight management application, which means that it does not require a highly performant connection to complete its work, and can serve as a central management layer for Delphix Engines globally.



6.1 Kubernetes

- [Installation and setup for Kubernetes](#) (see page 44)
- [Ingress setup](#) (see page 51)
- [Bootstrapping API keys](#) (see page 55)
- [DCT logs for Kubernetes](#) (see page 58)
- [Admin topics](#) (see page 58)

6.1.1 Installation and setup for Kubernetes

- ⓘ** Before getting started, Delphix recommends engaging your Kubernetes Admin for deployment guidance on available Kubernetes deployment platforms (AKS, EKS, etc.). Configurations like node sizing and persistent volume settings will need to be determined.

6.1.1.1 Hardware requirements


The hardware requirements for Data Control Tower (DCT) on Kubernetes are listed below. In addition to these requirements, inbound port 443 must be open for API clients, and outbound port 443 to engines. This is the minimum total resource request for the Kubernetes deployment of DCT. Individual service-level resource requests are contained in **values.yaml** file and can be overridden during deployment.

- **CPU:** 4-Core
- **CPU Architecture:** x86_64

- **Memory:** 16GB
- **Storage:** 50GB
- **Port:** 443

The recommended minimum 50 GB of storage is shared across the Kubernetes cluster (i.e. hosts). All pods and/or services use this storage for mounted volumes and other utilities, including image storage.

In a single node cluster, if shared volumes are not externalized, the host requires the full 50 GB of storage. If the persistent volume is mounted externally, the host requires 38 GB of storage, since the default storage required by the database (10 GB), gateway (1 GB) and masking (1 GB) draws from the external storage. The default storage configuration for the database and gateway can be modified in values.yaml.

 Many users may have default container settings as part of their Kubernetes or OpenShift infrastructure management. It is paramount to compare those default settings with the recommended minimum performance specifications (please engage your container infrastructure team to verify). If those default setting are lower, please update them to the minimum or higher.

For users who need to have limits set, you can start with the following sample configuration.

```
dataBookmarks:
  resources:
    requests:
      memory: "256Mi"
      cpu: "100m"
    limits:
      memory: "512Mi"
      cpu: "200m"
dataLibrary:
  resources:
    requests:
      memory: "256Mi"
      cpu: "100m"
    limits:
      memory: "512Mi"
      cpu: "200m"
database:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
    limits:
      memory: "1024Mi"
      cpu: "1500m"
gateway:
  resources:
    requests:
      memory: "512Mi"
```

```

    cpu: "1"
  limits:
    memory: "1024Mi"
    cpu: "1"
graphql:
  resources:
    requests:
      memory: "256Mi"
      cpu: "100m"
    limits:
      memory: "512Mi"
      cpu: "200m"
jobs:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
    limits:
      memory: "512Mi"
      cpu: "400m"
jobsCleaner:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
    limits:
      memory: "512Mi"
      cpu: "400m"
masking:
  resources:
    requests:
      memory: "512Mi"
      cpu: "500m"
    limits:
      memory: "1024Mi"
      cpu: "500m"

```

6.1.1.2 Kubernetes overview

Data Control Tower can be deployed in a matter of minutes, once a Kubernetes cluster has been identified and deployment details have been aligned with your Kubernetes administrator. The installation consists of three components:

- **Kubernetes cluster:** The identified infrastructure to which DCT will be deployed.
- **HELM:** This deploys DCT as a Kubernetes application by referencing HELM charts (.yaml files) that make up the DCT install either by an external helm repository (<https://dlpx-helm-dct.s3.amazonaws.com>¹¹, this is the quickest path to installing DCT, as it largely automated) or via

¹¹ <https://dlpx-helm-dct.s3.amazonaws.com/>

local install (this is accomplished by downloading the helm charts directly via the DCT .tar file on download.delphix.com¹²).

- `kubectl` : Is a command line tool that enables administrative communication with the deployed pods (most useful post-deployment or after an upgrade).

6.1.1.3 Installation requirements (Kubernetes)

DCT requires a running Kubernetes cluster; This could be an on-premises cluster, Azure AKS or AWS EKS cluster. DCT also requires a `kubectl` command line tool to interact with Kubernetes cluster and HELM for deployment on to the cluster.

Requirement	DCT recommended version	Comments
Kubernetes Cluster	1.25 or above	
HELM	3.9.0 or above	Install HELM as the package manager using the HELM installation ¹³ article.
kubectl	1.25.0 or above	HELM will internally refer to the kubeconfig file to connect to the Kubernetes cluster. The default kubeconfig file is present at location: <code>~/.kube/config</code> If the kubeconfig file needs to be overridden while running HELM commands, set the KUBECONFIG environment variable to the location of the kubeconfig file. To install kubectl follow the instructions at https://kubernetes.io/docs/tasks/tools/ .

6.1.1.4 Installing DCT

Add the DCT HELM repo with the following, which will link the DCT HELM repo to the local client HELM repo:

```
helm repo add dct-services https://dlpx-helm-dct.s3.amazonaws.com
```

Update the added repos with the following, which will update the local HELM charts metadata (index.yaml).

```
helm repo update
```

¹² <http://download.delphix.com>

¹³ <https://helm.sh/docs/intro/install/>

Pull the helm charts with the following, which is used to download the .tgz file.

```
helm pull dct-services/delphix-dct --version x.0.0
```

The downloaded file is then extracted using the following command (where `x.0.0` should be changed to the version of DCT being installed):

```
tar -xvf delphix-dct-x.0.0.tgz
```

Update the following properties values.yaml which is present in the extracted folder `delphix-dct`.

- To generate the bootstrap APIKey, set `apiKeyCreate: true`.
- Provide image credentials to pull images from docker registry.
 - `username: <retrieved from download.delphix.com>`
 - `password: <retrieved from download.delphix.com>`

For getting the image credentials, visit the [Delphix DCT Download](#)¹⁴ page and login with your customer login credentials. Once logged in, select the **DCT Helm Repository** link and accept the Terms and Conditions. Use password from the popup screen, as shown below.

Helm Credentials: DCT Helm Repository ✕

These credentials are linked to your organization, rather than your individual Delphix Support account, and will not expire, unless unused for an extended period of time.

Helm URL	<code>https://dlpx-helm-dct.s3.amazonaws.com/</code>	
Username	<code>BDOAXpNq-y4M</code>	
Password	<code>01e89511-f200-40a0-8680-96a258cf64f3</code>	
Organization	Delphix	

Close

¹⁴ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

OFFLINE MODE

If you do not have direct access to Delphix HELM and Docker repositories, download the HELM charts and Docker images packages as [delphix-dct-x.0.0.tar.gz](#)¹⁵, which is uploaded on the [download site](#)¹⁶.

- Extract the downloaded .tgz file to get the HELM chart and Docker images.
- Load extracted Docker images using following:
 - `for image in *tar; do sudo docker load --input $image; done`
- RE-tag and push Docker images in your local repository, tag images in the following format:
 - e.g., `docker tag registry.delphix.com/delphix-dct:app-x.0.0 <local registry url>:app-x.0.0`
 - **Note:** Re-tag all remaining DCT images in the same format.
- Update values.yaml for registry name and image credentials:
 - `registry: <local registry url>`
 - `username: <local registry username>`
 - `password: <local registry password>`
- Deploy.

In version 8.0.0 and above, instead of username/password, an option is available to use credentials from a pre-existing [Kubernetes Secret](#)¹⁷. To do so, instead of providing a username/password, users must create the Kubernetes Secret in the same namespace as the one used for DCT, and reference the registryKey as follows:

```
imageCredentials:
  # registry to pull docker images from.
  registry: dct.download.delphix.com/delphix-dct
  # username to login to docker registry. Do not set if registryKey is set.
  username:
  # password to login to docker registry. Do not set if registryKey is set.
  password:
  # Name of an existing docker registry key to use to pull images.
  registryKey: <insert-secret-name-here>
```

After updating the values.yaml, install it using the following command:

¹⁵ <https://download.delphix.com/file/12747/delphix-dct-9.0.0.tar.gz>

¹⁶ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

¹⁷ <https://kubernetes.io/docs/concepts/configuration/secret/>

```
helm install dct-services delphix-dct
```

delphix-dct is the name of the folder which was extracted in the previous step. **dct-services** is the chart name which is given for this deployment.

In the above directory structure, the **values.yaml** file contains the properties which are configured above. If you want to configure additional properties in values.yaml, create a values.yaml using the sample below. Deploy DCT using following command to use the custom values.yaml:

```
helm install dct-services -f <path to edited values.yaml> <directory path of the extracted chart>
```

6.1.1.4.1 Sample values.yaml file

A sample values.yaml file can be downloaded below.

values.yaml

[\(see page 44\)](#)

Once deployment is complete, check the status of the deployment using the following command (where **X.0.0** should be changed to the version of DCT being installed):

```
helm list
NAME          NAMESPACE    REVISION    UPDATED
STATUS      CHART          APP VERSION
dct-services default       1           2023-01-10 19:33:41.713202 -0900
deployed    delphix-dct-x.0.0 x.0.0
```

- Assuming an ingress controller configuration on the Kubernetes cluster is present, when accessing DCT after the deployment, the ingress controller rule needs to be added for proxy service, along with port 443 (if SSL is enabled) and port 80 (if SSL is disabled).

6.1.2 Ingress setup

Ingress exposes HTTP and HTTPS routes from outside the cluster to DCT running within the cluster. Find out more about Ingress from their [official documentation](#)¹⁸.

- The exact steps to setup an Ingress vary by Kubernetes vendor and company policies. This section provides non-exhaustive instructions for a basic setup, but please ask your Kubernetes cluster administrator for guidance.

The proxy pod (which comes with DCT) runs an Nginx HTTP server which must be the only target of the Ingress rules, redirecting all external traffic to it. Out of the box, the pod accepts requests over HTTPS on port 443, using a self-signed certificate.

6.1.2.1 Expose proxy HTTP port (80) for non-encrypted traffic

After setting up an Ingress, TLS will be terminated by the HTTP server/load balancer/proxy implementing the Ingress, and not DCT. First, disable the TLS (SSL) configuration of DCT itself, making it expose port 80 for non encrypted traffic. To do so, edit the **values.yaml** to unset the `useSSL` property.

Either expose proxy on SSL port or non SSL port:

```
useSSL: false
```

Then run `helm upgrade` to apply the changes:

```
helm upgrade dct-services -f <path to edited values.yaml> <directory path of the extracted chart>
```

Now, the proxy pod accepts unencrypted traffic on port 80.

¹⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.27/#ingress-v1-networking-k8s-io>

6.1.2.2 Ingress controller installation and route creation

An [Ingress controller](#)¹⁹ is required to continue. Expand a section below based on your Kubernetes environment to show the corresponding **Ingress controller installation** and **Ingress route creation** instructions.

Microsoft Azure AKS

6.1.2.2.1 Ingress controller installation

Please follow these [instructions](#)²⁰ to install an Nginx Ingress controller. A simple setup can be installed with these commands:

```

NAMESPACE=ingress-basic
helm repo add ingress-nginx https://kubernetes.github.io/ingress-nginx
helm repo update
helm install ingress-nginx ingress-nginx/ingress-nginx \
  --create-namespace \
  --namespace $NAMESPACE \
  --set controller.service.annotations."service\.beta\.kubernetes\.io/azure-load-balancer-health-probe-request-path"/healthz

```

6.1.2.2.2 Ingress route creation

Create a file named **ingress.yaml**.

```

apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: dct-ingress
  annotations:
    nginx.ingress.kubernetes.io/ssl-redirect: "true"
spec:
  ingressClassName: nginx
  rules:
    - http:
        paths:
          - path: /
            pathType: Prefix
            backend:
              service:
                name: proxy
                port:
                  number: 80

```

¹⁹ <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers>

²⁰ <https://learn.microsoft.com/en-us/azure/aks/ingress-basic>

Apply the Ingress resource with `kubectl apply` :

```
kubectl apply -f ingress.yaml --namespace=ingress-basic
```

To configure TLS, see [Use TLS with an Ingress controller](#)²¹.

Amazon AWS EKS

6.1.2.2.3 Ingress controller installation

Please follow these [instructions](#)²² to install an [AWS load balancer controller](#)²³ (An Ingress controller which configures AWS application load balancers).

6.1.2.2.4 Ingress route creation

Create a file named **ingress.yaml**, replacing the value of `certificate-arn` in the example below with the ARN of the certificate you want to use for the HTTPs endpoint.

```
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: dct-ingress
  annotations:
    kubernetes.io/ingress.class: alb
    alb.ingress.kubernetes.io/scheme: internal
    alb.ingress.kubernetes.io/target-type: ip
    alb.ingress.kubernetes.io/ssl-redirect: '443'
    alb.ingress.kubernetes.io/listen-ports: '[{"HTTP": 80}, {"HTTPS":443}]'
    alb.ingress.kubernetes.io/certificate-arn: arn:aws:acm:us-west-2:xxxxx:certificat
e/xxxxxxx
spec:
  rules:
    - http:
        paths:
          - path: /
            pathType: Prefix
            backend:
              service:
                name: proxy
                port:
                  number: 80
```

²¹ <https://learn.microsoft.com/en-us/azure/aks/ingress-tls>

²² <https://docs.aws.amazon.com/eks/latest/userguide/alb-ingress.html>

²³ <https://docs.aws.amazon.com/eks/latest/userguide/aws-load-balancer-controller.html>

Alternatively, you may use [certificate discovery](#)²⁴ to have the ALB select a matching certificate from [AWS Certificate manager](#)²⁵ based on the host name.

```
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: dct-ingress
  annotations:
    kubernetes.io/ingress.class: alb
    alb.ingress.kubernetes.io/scheme: internal
    alb.ingress.kubernetes.io/target-type: ip
    alb.ingress.kubernetes.io/ssl-redirect: '443'
    alb.ingress.kubernetes.io/listen-ports: '[{"HTTP": 80}, {"HTTPS":443}]'
spec:
  tls:
  - hosts:
    - www.example.com
  rules:
  - http:
    paths:
    - path: /
      pathType: Prefix
      backend:
        service:
          name: proxy
          port:
            number: 80
```

Apply the Ingress resource with `kubectl apply` :

```
kubectl apply -f ingress.yaml --namespace=ingress-basic
```

This creates an [application load balancer](#)²⁶, which forwards all traffic to DCT.

Other

6.1.2.2.5 Ingress controller installation

For **self-hosted Kubernetes**, you can install an Nginx Ingress controller with:

```
helm upgrade --install ingress-nginx ingress-nginx --repo https://
kubernetes.github.io/ingress-nginx --namespace ingress-nginx --create-namespace
```

²⁴ https://kubernetes-sigs.github.io/aws-load-balancer-controller/v2.5/guide/ingress/cert_discovery/

²⁵ <https://aws.amazon.com/certificate-manager/>

²⁶ <https://docs.aws.amazon.com/eks/latest/userguide/alb-ingress.html>

Reference [environment specific instructions](#)²⁷ for a complete list of supported platforms and specific instructions.

6.1.2.2.6 Ingress route creation

Create a file name **ingress.yaml**.

```
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: dct-ingress
  annotations:
    nginx.ingress.kubernetes.io/ssl-redirect: "true"
spec:
  ingressClassName: nginx
  rules:
    - http:
        paths:
          - path: /
            pathType: Prefix
            backend:
              service:
                name: proxy
                port:
                  number: 80
```


Apply the Ingress resource with `kubectl apply` :

```
kubectl apply -f ingress.yaml --namespace=ingress-basic
```

Review the Ingress-Nginx [instructions](#)²⁸ to setup TLS.

6.1.3 Bootstrapping API keys

API keys are the default method to authenticate with DCT. This is done by including the key in the [HTTP Authorization request header](#)²⁹ with type `apk` .

 API keys are long-live tokens and as a result, do not automatically expire in the future. They remain valid until they are deleted or destroyed from DCT.

²⁷ <https://kubernetes.github.io/ingress-nginx/deploy/>

²⁸ <https://kubernetes.github.io/ingress-nginx/user-guide/tls/>


²⁹ <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Authorization>

A cURL example using an example key of

`1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3` would appear as:

```
curl --header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3'
```

cURL (like web browsers and other HTTP clients) will not connect to DCT over HTTPS unless a valid TLS certificate has been configured for the Nginx server. If this [configuration step](#) (see page 116) has not been performed yet and the risk is comprehended, you may disable the check in the HTTP client. For instance, this can be done with cURL using the `--insecure` flag.

 The cURL version must be 7.43 or higher.

6.1.3.1 Bootstrap first API key

There is a special process to bootstrap the creation of the first API key. This first API key should only be used to create another key and then promptly deleted, since the bootstrap API will appear in the logs. This process can be repeated as many times as needed, for example, in a case where existing API keys are lost or have been deleted.

Once the application is started, edit the **values.yaml** file and modify the following lines, to set the `apiKeyCreate` to the string value `true`. Toggle this value to create/seed bootstrap API key:

```
apiKeyCreate: true
```

Upgrade DCT with:

```
helm upgrade dct-services <directory path of the extracted chart>
```

If the values.yaml file needs to be overridden from outside, then use:

```
helm upgrade dct-services -f <path_to_values.yaml> <directory path of the extracted chart>
```

You will see the following output in the logs for the *gateway* pod (the key will be different from this example):

```
NEWLY GENERATED API KEY: 1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3
```

Logs for a gateway pod can be accessed using:


```
kubectl logs <gateway-pod-name> -n dct-services
```

`gateway-pod-name` will be of the format `gateway-xxx` and can be found using the following command:

```
kubectl get pods -n dct-services
```

Copy the API Key, it can now be used to authenticate with DCT. Remember that the API Key value must be prefixed with **APK**. An example cURL command with the above API Key appears as follows:

```
curl --header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3'
```

Edit the **values.yaml** file to set the `apiKeyCreate` environment variable value back to `false` and upgrade DCT again with:

```
helm upgrade dct-services <directory path of the extracted chart>
```

If the values.yaml file needs to be overridden from outside, then use:

```
helm upgrade dct-services -f <path_to_values.yaml> <directory path of the extracted chart>
```

6.1.3.1.1 Create and manage API Keys

The initial API key created should be used to create a new admin secure key. This is done by creating a new *Account* entity and setting the `generate_api_key`. The "username" attribute should be the desired name to uniquely identify the account.

 If the cURL version is below 7.43, replace `--data-raw` option with `--data`.

```
curl --location --request POST 'https://<hostname>/v2/management/accounts' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3' \
--data-raw '{
  "username": "secure-key",
  "generate_api_key": true,
  "is_admin": true
```

```
}'
```

A response should be received similar to the lines below:

```
{
  "id": 2,
  "token": "2.vCfC0MnpySYZLshuxap2aZ7xqBKAnQvV7hFnobe7xuNlHS9AF2NqnV9XXw4UyET6"
  "username": "secure-key"
}
```

Now that the new and secure API key is created, the old one must be deleted for security reasons since the key appeared in the logs. To do this make the following request:

```
curl --location --request DELETE 'https://<hostname>/v2/management/api-clients/<id>' \
  --header 'Content-Type: application/json' \
  --header 'Accept: application/json' \
  --header 'Authorization: apk
2.vCfC0MnpySYZLshuxap2aZ7xqBKAnQvV7hFnobe7xuNlHS9AF2NqnV9XXw4UyET6'
```

The `id` referenced above is the numeric id of the Account. It is the integer before the period in the token.

For example, the id of

```
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3
```

is 1.

Finally, to list all of the current Accounts, make the following request:

```
curl --location --request GET 'https://<hostname>/v2/management/accounts/' \
  --header 'Content-Type: application/json' \
  --header 'Accept: application/json' \
  --header 'Authorization: apk <your API key>'
```

6.1.4 DCT logs for Kubernetes

All DCT containers log to stdout and stderr so that their logs are processed by Kubernetes. To view container level logs running on the Kubernetes cluster use:

```
kubectl logs <pod_name> -n dct-services
```

Log aggregators can be configured to read from `stdout` and `stderr` for all of the pods as per the requirements.

6.1.5 Admin topics

- [Deployment upgrade for Kubernetes \(see page 59\)](#)

- [Factory reset DCT for Kubernetes \(see page 61\)](#)

6.1.5.1 Deployment upgrade for Kubernetes

This page covers the upgrade process for DCT deployments on Kubernetes.

- Starting from DCT 12.0.0 we have introduced one new service(container) named Masking. Memory and CPU requirement for this newly added service(container) is 512Mi and 500m respectively.

This container creates a new persistent volume of 1GB dynamically, if you are using pre-existing static persistent volumes for other existing pods(i.e. gateway and database) then similar to these pods, please create a new persistent volume claim for masking pod too and update the `maskingPvcName` property in the values.yaml with the name of static persistent volume claim.

For customising masking service resources(request and limits) and properties please refer values.yaml file in the helm chart of current release.

Create a new folder called **dct-x.0.0**, where `x.0.0` should be changed to the version of DCT being installed (e.g. if on 5.0.2, it would be 6.0.0).

```
mkdir dct-x.0.0
```

Update the added repos with the following, which will update the local HELM charts metadata (index.yaml).

```
helm repo update
```

Pull the HELM charts with the following:

- This command will download a file named **delphix-dct-x.0.0.tgz** in the folder `dct-x.0.0`.

```
cd dct-x.0.0  
helm pull dct-services/delphix-dct --version x.0.0
```

The downloaded file is then extracted using the following command (where `x.0.0` should be changed to the version of DCT being installed):

```
tar -xvf delphix-dct-x.0.0.tgz
```

Which will extract into the following directory structure:

```
delphix-dct
|- values.yaml
|- README.md
|- Chart.yaml
|- templates
   |-<all templates files>
```

Create a new values.yaml inside to the dct-x.0.0 folder (e.g. parallel to delphix-dct folder), use sample values.yaml given below. Copy the configured properties from previous version to this newly created values.yaml.

values.yaml

(see page 59)



This values.yaml file should only contain the modified values from the previous version of deployment and not the entire values.yaml file as it is.

Update the username and password in values.yaml. It can be obtained from <https://download.delphix.com>³⁰. Here are some notes in regards to this step in the process:

- This username and password update in values.yaml is only required if the user using DCT provided a Docker Registry directly in the deployment (i.e. values.yaml).

³⁰ <https://download.delphix.com/>

- Username and password update is required only if username AWS and temporary password is used.

OFFLINE MODE

If you do not have direct access to Delphix HELM and Docker repositories, download the HELM charts and Docker images packages as [delphix-dct-x.0.0.tar.gz](#)³¹, which is uploaded on the [download site](#)³².

- Extract the downloaded .tgz file to get the HELM chart and Docker images.
- Load extracted Docker images using following:
 - for image in *tar; do sudo docker load --input \$image; done
- RE-tag and push Docker images in your local repository, tag images in the following format:
 - e.g., docker tag registry.delphix.com/delphix-dct:app-x.0.0 <local registry url>:app-x.0.0
 - Note:** Re-tag all remaining DCT images in the same format.
- Update values.yaml for registry name and image credentials:
 - registry: <local registry url>
 - username: <local registry username>
 - password: <local registry password>

- Run the HELM upgrade command:

```
helm upgrade -f values.yaml dct-services delphix-dct
```

6.1.5.2 Factory reset DCT for Kubernetes

To clean DCT installation run following command:

```
helm delete dct-services
```

³¹ <https://download.delphix.com/file/12747/delphix-dct-9.0.0.tar.gz>

³² <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>



This process will delete services pod and database both.

6.2 OpenShift

- [Installation and setup for OpenShift](#) (see page 62)
- [OpenShift authentication](#) (see page 71)
- [DCT logs for OpenShift](#) (see page 72)
- [Admin topics for OpenShift](#) (see page 72)

6.2.1 Installation and setup for OpenShift

6.2.1.1 Hardware requirements

The hardware requirements for Data Control Tower to deploy on OCP are listed below. In addition to these requirements, inbound port 443 or 80 must be open for API clients. This is the minimum total resource requirement for the deployment.

- **CPU:** 4-Core
- **CPU Architecture:** x86_64
- **Memory:** 16GB
- **Storage:** 50GB
- **Port:** 443

The recommended minimum 50 GB of storage is shared across the Kubernetes cluster (i.e. hosts). All pods and/or services use this storage for mounted volumes and other utilities, including image storage.

In a single node cluster, if shared volumes are not externalized, the host requires the full 50 GB of storage. If the persistent volume is mounted externally, the host requires 38 GB of storage, since the default storage required by the database (10 GB), gateway (1 GB) and masking (1 GB) draws from the external storage. The default storage configuration for the database and gateway can be modified in values.yaml.



Many users may have default container settings as part of their Kubernetes or OpenShift infrastructure management. It is paramount to compare those default settings with the recommended minimum performance specifications (please engage your container infrastructure team to verify). If those default setting are lower, please update them to the minimum or higher.

For users who need to have limits set, you can start with the following sample configuration.

```
dataBookmarks:
```

```
resources:
  requests:
    memory: "256Mi"
    cpu: "100m"
  limits:
    memory: "512Mi"
    cpu: "200m"
dataLibrary:
  resources:
    requests:
      memory: "256Mi"
      cpu: "100m"
    limits:
      memory: "512Mi"
      cpu: "200m"
database:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
    limits:
      memory: "1024Mi"
      cpu: "1500m"
gateway:
  resources:
    requests:
      memory: "512Mi"
      cpu: "1"
    limits:
      memory: "1024Mi"
      cpu: "1"
graphql:
  resources:
    requests:
      memory: "256Mi"
      cpu: "100m"
    limits:
      memory: "512Mi"
      cpu: "200m"
jobs:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
    limits:
      memory: "512Mi"
      cpu: "400m"
jobsCleaner:
  resources:
    requests:
      memory: "256Mi"
      cpu: "200m"
```

```

limits:
  memory: "512Mi"
  cpu: "400m"
masking:
resources:
  requests:
    memory: "512Mi"
    cpu: "500m"
  limits:
    memory: "1024Mi"
    cpu: "500m"

```

6.2.1.2 Installation requirements (OpenShift)

DCT requires a running OpenShift cluster to run, `oc` command line tool to interact with OpenShift cluster, and HELM for deployment on to the cluster.

Requirement	DCT Recommended Version	Comments
OpenShift Cluster	4.12 or above	
HELM	3.9.0 or above	HELM installation should support HELM v3. More information on HELM can be found at https://helm.sh/docs/ . To install HELM, follow the installation instructions at https://helm.sh/docs/intro/install/ .
oc	4.11.3 or above	To install oc follow the instructions at https://docs.openshift.com/container-platform/4.8/cli_reference/openshift_cli/getting-started-cli.html .



If an intermediate HELM repository is to be used instead of the default Delphix HELM repository, then the repository URL, username, and password to access this repository needs to be configured in the `values.yaml` file under `imageCredentials` section.

6.2.1.3 Installation process

6.2.1.3.1 Jumpbox setup

6.2.1.3.1.1 OC login

Run the OC login command to authenticate OpenShift CLI with the server:

```
oc login https://openshift1.example.com --token=<<token>>
```

6.2.1.3.1.2 Verify KubeConfig

HELM will use the configuration file inside the **\$HOME/.kube/** folder to deploy artifacts on an OpenShift cluster.

Be sure the config file has the cluster context added, and the current-context is set to use this cluster. To verify the context, run this command:

```
oc config current-context
```

6.2.1.3.1.3 Create a new project

Create a new project named **dct-services** using the command below:

```
oc new-project dct-services --description="DCT Deployment project" --display-name="dct-services"
```

6.2.1.3.1.4 Installing Helm

Install HELM using the following installation instructions mentioned at <https://helm.sh/docs/intro/install/>.

6.2.1.3.2 Deploy DCT chart

Add the DCT HELM repo with the following, which will link the DCT HELM repo to the local HELM repo:

```
helm repo add dct-services https://d1px-helm-dct.s3.amazonaws.com
```

Update the added repos with the following, which will update the local HELM charts metadata (index.yaml):

```
helm repo update
```

Pull the HELM charts with the following, which is used to download the .tgz file:

```
helm pull dct-services/delphix-dct --version x.0.0
```

The downloaded file is then extracted using the following command (where `x.0.0` should be changed to the version of DCT being installed):

```
tar -xvf delphix-dct-x.0.0.tgz
```

Update the following properties values.yaml which is present in the extracted folder `delphix-dct`.

- To generate the bootstrap APIKey, set `apiKeyCreate: true`.
- Provide image credentials to pull images from docker registry.
 - `username: <retrieved from download.delphix.com>`
 - `password: <retrieved from download.delphix.com>`
- `isOpenshift` should be set to `true`.

For getting the image credentials, visit the [Delphix DCT Download](#)³³ page and login with your user login credentials. Once logged in, select the **DCT Helm Repository** link and accept the Terms and Conditions. Use password from the popup screen, as shown below.

Helm Credentials: DCT Helm Repository ✕

These credentials are linked to your organization, rather than your individual Delphix Support account, and will not expire, unless unused for an extended period of time.

Helm URL	<code>https://dlpx-helm-dct.s3.amazonaws.com/</code>	
Username	<code>BDOAXpNq-y4M</code>	
Password	<code>01e89511-f200-40a0-8680-96a258cf64f3</code>	
Organization	Delphix	

Close

³³ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

OFFLINE MODE

If you do not have direct access to Delphix HELM and Docker repositories, download the HELM charts and Docker image packages as [delphix-dct-x.0.0.tar.gz](#)³⁴, which is uploaded on the [download site](#)³⁵.

- Extract the downloaded .tgz file to get the HELM chart and Docker images.
- Load extracted Docker images using the following command:
 - For an image in *tar, use `sudo docker load --input $image; done`.
- RE-tag and push Docker images in your local repository, tag images in the following format:
 - e.g. `docker tag registry.delphix.com/delphix-dct:app-x.0.0 <local registry url>:app-x.0.0`.
 - **Note:** Re-tag all remaining DCT images in the same format:
- Update values.yaml for registry name and image credentials:
 - `registry: <local registry url>`
 - `username: <local registry username>`
 - `password: <local registry password>`
- Deploy.

6.2.1.3.2.1 Find and update fsGroup in values.yaml file

The **fsGroup** field is used to specify a supplementary group ID. All processes of the container, the owner of the volume, and any files created on the volume are also part of this supplementary group ID.

For OpenShift deployment, this value need to be specified in the values.yaml file.

Find the allowed supplementary group range:

```
oc get project dct-services -o yaml
```

A response should appear as follows:

```
apiVersion: project.openshift.io/v1
kind: Project
metadata:
```

³⁴ <https://download.delphix.com/file/12747/delphix-dct-9.0.0.tar.gz>

³⁵ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

```

annotations:
  openshift.io/description: ""
  openshift.io/display-name: ""
  openshift.io/requester: cluster-admin
  openshift.io/sa.scc.mcs: s0:c32,c4
  openshift.io/sa.scc.supplemental-groups: 1001000000/10000
  openshift.io/sa.scc.uid-range: 1001000000/10000
creationTimestamp: "2023-01-18T10:33:04Z"
labels:
  kubernetes.io/metadata.name: dct-services
  pod-security.kubernetes.io/audit: restricted
  pod-security.kubernetes.io/audit-version: v1.24
  pod-security.kubernetes.io/warn: restricted
  pod-security.kubernetes.io/warn-version: v1.24
name: dct-services
resourceVersion: "99974"
uid: ccdd5c9f-2ce5-49b4-91a7-662e0598b63b
spec:
  finalizers:
  - kubernetes
status:
  phase: Active

```

Copy the first value from the `openshift.io/sa.scc.supplemental-groups` line, before the slash (e.g. 1001000000). Paste this value in the `values.yaml` file:

```

# Define SecurityContextConstraints for the pod
podSecurityContext:
  fsGroup: 1001000000

```

6.2.1.3.2.2 Deploy DCT

Run the following command to deploy the DCT chart (where `x.0.0` should be changed to the version of DCT being installed):

```
helm install dct-services delphix-dct
```

delphix-dct is the name of the folder which was extracted in the previous step. `dct-services` is the chart name which is given for this deployment.

In the above directory structure, the **values.yaml** file contains the properties which are configured above. Deploy DCT using following command to use the custom `values.yaml`:

```
helm install dct-services -f <path to edited values.yaml> <directory path of the extracted chart>
```

6.2.1.3.2.3 Verify deployment

All the images will be downloaded and then deployed. If some pods restarted at the startup, this is expected. After some time, a total of 9 pods will be in running status and one job pod will be in completed status.

```
oc get pods -n dct-services
```

6.2.1.3.2.4 Find API key

For the very first deployment bootstrap API key will be printed in logs, please view gateway pod logs and find for “NEWLY GENERATED API KEY”. the value is the API key.

```
oc logs <gateway-pod-name> -n dct-services
```

6.2.1.4 Configure Ingress

DCT only works with HTTPS Ingress, the UI does not support HTTP.

6.2.1.4.1 Creating route

To create a route, you can use the OpenShift console and create a new one for the DCT service.

If SSL is terminated at this route, only then should the useSSL value in values.yaml be updated to false, so that 80 port will be exposed in proxy service and can be used to configure the route. The following screenshot shows the route that forwards requests to 80 port of proxy service:

Project: apigw-services

Configure via: Form view YAML view

Name *
dct
A unique name for the Route within the project.

Hostname
dct.delphix.com
Public hostname for the Route. If not specified, a hostname is generated.

Path
/
Path that the router watches to route traffic to the service.

Service *
proxy
Service to route to.
[Add alternate Service](#)

Target port *
80 → 8083 (TCP)
Target port for traffic.

Security
 Secure Route
Routes can be secured using several TLS termination types for serving certificates.

TLS termination *
Edge

Insecure traffic
Redirect
Policy for traffic on insecure schemes like HTTP.

Certificates
TLS certificates for edge and re-encrypt termination. If not specified, the router's default certificate is used.

If SSL is not terminated at the Route level, then create a PassThrough route and use 443 port of the proxy service, and configure the SSL certificate and key in the values.yaml file:

Project: apigw-services

Routing is a way to make your application publicly visible.

Configure via: Form view YAML view

Name *
dct
A unique name for the Route within the project.

Hostname
dct.delphix.com
Public hostname for the Route. If not specified, a hostname is generated.

Path
/
Path that the router watches to route traffic to the service.

Service *
proxy
Service to route to.
[Add alternate Service](#)

Target port *
443 → 8443 (TCP)
Target port for traffic.

Security
 Secure Route
Routes can be secured using several TLS termination types for serving certificates.

TLS termination *
Passthrough

Insecure traffic
Redirect
Policy for traffic on insecure schemes like HTTP.

[Create](#) [Cancel](#)

6.2.2 OpenShift authentication

6.2.2.1 Introduction

DCT uses Nginx/OpenResty as an HTTP server and a reverse proxy for the application. Using the default configuration, all connections to DCT are over HTTPS and require the user to authenticate. There are three supported methods for authentication; API keys, Username/Password, and OpenID Connect.

6.2.2.2 Enable OAuth2 authentication

By default APIKey authentication will be enabled and when DCT starts it will generate a new [API key \(see page 98\)](#) in logs if you want to enable openid connect authentication then follow below procedure:

Update the below properties in the **values.yaml** file and restart DCT:

```
# flag to enable api_key based authentication
apiKeyEnabled: false
# flag to enable OAuth2 based authentication
openIdEnabled: true
# URL of the discovery endpoint as defined by the OpenId Connect Discovery
specification. This needs to be set if 'openIdEnabled' is set to true
openIdServerUrl: https://delphix.okta.com/oauth2/default/.well-known/oauth-
authorization-server
# OAuth2 jwt claim name that should be used as client_id
jwtClaimForClientId: sub
# OAuth2 jwt claim name that should be used as client_name
jwtClaimForClientName: sub
```

6.2.3 DCT logs for OpenShift

All DCT containers log to **stdout** and **stderr**, so that their logs are processed by OpenShift. To view container level logs running on the OpenShift cluster, use this command:

```
oc logs <pod_name> -n dct-services
```


Log aggregators can be configured to read from **stdout** and **stderr** for all of the pods as per the requirements.

6.2.4 Admin topics for OpenShift

- [Deployment upgrade for OpenShift \(see page 72\)](#)
- [Factory reset DCT for OpenShift \(see page 75\)](#)

6.2.4.1 Deployment upgrade for OpenShift

This page covers the upgrade process for DCT deployments on Kubernetes.

-  Starting from DCT 12.0.0 we have introduced one new service(container) named Masking. Memory and CPU requirement for this newly added service(container) is 512Mi and 500m respectively.

This container creates a new persistent volume of 1GB dynamically, if you are using pre-existing static persistent volumes for other existing pods(i.e. gateway and database) then similar to these pods, please create a new persistent volume claim for masking pod too and update the `maskingPvcName` property in the values.yaml with the name of static persistent volume claim.

For customising masking service resources(request and limits) and properties please refer values.yaml file in the helm chart of current release.

Create a new folder called **dct-x.0.0** where `x.0.0` should be changed to the version of DCT being installed (e.g. if on 5.0.2, it would be 6.0.0).

```
mkdir dct-x.0.0
```

Update the added repos with the following, which will update the local HELM charts metadata (index.yaml):

```
helm repo update
```

Pull the HELM charts:



This command will download a file named **delphix-dct-x.0.0.tgz** in the folder `dct-x.0.0`.

```
cd dct-x.0.0
helm pull dct-services/delphix-dct --version x.0.0
```

The downloaded file is then extracted using the following command (where `x.0.0` should be changed to the version of DCT being installed):

```
tar -xvf delphix-dct-x.0.0.tgz
```

Which will extract into the following directory structure:

```
delphix-dct
|- values.yaml
|- README.md
|- Chart.yaml
|- templates
  |-<all templates files>
```

Create a new values.yaml inside to the `dct-x.0.0` folder (e.g. parallel to `delphix-dct` folder), use sample values.yaml given below. Copy the configured properties from previous version to this newly created values.yaml.

values.yaml

(see page 72)

[-] The values.yaml file contains only modified values from the previous version of deployment.

Update the username and password in values.yaml. It can be obtained from <https://download.delphix.com>³⁶. Here are some notes in regards to this step in the process:

- This username and password update in values.yaml is only required if the user using Delphix provided a Docker Registry directly in the deployment (i.e. values.yaml).

[-] Username and password update in values.yaml is required only if username AWS and temporary password(12 hours expiry) is used.

[-] OFFLINE MODE

If you do not have direct access to Delphix HELM and Docker repositories, download the HELM charts and Docker images packages as [delphix-dct-x.0.0.tar.gz](#)³⁷, which is uploaded on the [download site](#)³⁸.

- Extract the downloaded .tgz file to get the HELM chart and Docker images.
- Load extracted Docker images using following:
 - `for image in *tar; do sudo docker load --input $image; done`

³⁶ <https://download.delphix.com/>

³⁷ <https://download.delphix.com/file/12747/delphix-dct-9.0.0.tar.gz>

³⁸ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

- RE-tag and push Docker images in your local repository, tag images in the following format:
 - e.g., `docker tag registry.delphix.com/delphix-dct:app-x.0.0 <local registry url>:app-x.0.0`
 - **Note:** Re-tag all remaining DCT images in the same format.
- Update values.yaml for registry name and image credentials:
 - `registry: <local registry url>`
 - `username: <local registry username>`
 - `password: <local registry password>`
- Deploy.


- Run the HELM upgrade command:

```
helm upgrade -f values.yaml dct-services delphix-dct
```


6.2.4.2 Factory reset DCT for OpenShift

To clean DCT installation run following command:

```
helm delete dct-services
```

 This process will delete both services pod and database.


6.3 Docker Compose

 Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)³⁹ or [OpenShift](#)⁴⁰.

- [Installation and setup for Docker Compose](#) (see page 76)
- [Bootstrapping API Keys](#) (see page 79)
- [Custom configuration](#) (see page 81)
- [Docker logs](#) (see page 83)
- [Migration topics](#) (see page 83)
- [Admin topics for Docker Compose](#) (see page 90)

6.3.1 Installation and setup for Docker Compose

 Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁴¹ or [OpenShift](#)⁴².

6.3.1.1 Hardware requirements

The hardware requirements for Data Control Tower are listed below. In addition to these requirements, inbound port 443 must be open for API clients, and outbound port 443 to engines.

- **CPU:** 4-Core
- **CPU Architecture:** x86_64
- **Memory:** 16GB
- **Storage:** 50GB
- **Port:** 443

³⁹ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁴⁰ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

⁴¹ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁴² <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

6.3.1.2 Installation requirements (Docker Compose)

DCT **requires** Docker and Docker Compose to run, thus, Linux versions and distributions that have been verified to work with Docker are supported. To see a list of supported distributions, please reference this [Docker article](#)⁴³.

This example uses a [Docker installation](#)⁴⁴ and is completed on an Ubuntu 20.04 VM.

To begin, uninstall any old versions of Docker.

```
sudo apt-get remove docker docker-engine docker.io containerd runc
```

Next, update the package lists and install Docker.

```
sudo apt-get update
sudo apt-get install docker.io
```

Last, [install Docker Compose](#)⁴⁵.

```
sudo curl -L "https://github.com/docker/compose/releases/download/1.29.1/docker-
compose-$(uname -s)-$(uname -m)" -o /usr/local/bin/docker-compose
sudo chmod +x /usr/local/bin/docker-compose
```



Docker-Compose is packaged with Docker engine version 20.10.15 and up.

6.3.1.2.1 Running Docker as non-root (optional)

To avoid prefacing the Docker command with `sudo`, create a Unix group called `docker` and add users to it. When the Docker daemon starts, it creates a Unix socket accessible by members of the Docker group. See [Docker Post Installation](#)⁴⁶ documentation for details.

```
sudo groupadd docker
sudo usermod -aG docker $USER
```

43 <https://docs.docker.com/engine/install/#server>

44 <https://docs.docker.com/engine/install/>

45 <https://docs.docker.com/compose/install/>

46 <https://docs.docker.com/engine/install/linux-postinstall/>

6.3.1.3 Unpack and install DCT

Once Docker and Docker Compose are installed, DCT can be installed. Begin by downloading the latest version of the tarball from the [Delphix Download site](#)⁴⁷. Next, transfer the file to the Linux machine where Docker is installed. Run the following commands to extract the containers and load them into Docker (where `x.0.0` should be changed to the version of DCT being installed):

```
tar -xzf delphix-dct-x.0.0.tar.gz
for image in *tar; do sudo docker load --input $image; done
```

6.3.1.4 Run DCT

To run DCT, navigate to the location of the extracted **docker-compose.yaml** file from the tarball and run the following command. Using `-d` in the command will start up the application in the background.

```
sudo docker-compose up -d
```

Running `docker ps` should show 10 containers up and running:

```
sudo docker ps
CONTAINER ID   IMAGE                                COMMAND                                CREATED
STATUS        PORTS                NAMES                                COMMAND                                CREATED
75a9df0cae07  delphix-dct-proxy:x.0.0            "/sbin/tini -- /boot..."           7 seconds
ago          Up 4 seconds        0.0.0.0:443->8443/tcp delphix-dct-proxy:x.0.0            7 seconds
a23f4f0e220   delphix-dct-app:x.0.0              "java -jar /opt/delp..."           7 seconds
ago          Up 5 seconds        delphix-dct-app:x.0.0
96ba8018fa03  delphix-dct-data-library:x.0.0     "/usr/bin/tini -- ./..."           7 seconds
ago          Up 5 seconds        delphix-dct-data-library:x.0.0
8e5b1e671acc  delphix-dct-jobs:x.0.0             "/usr/bin/tini -- ./..."           7 seconds
ago          Up 5 seconds        delphix-dct-jobs:x.0.0
96049058f025  delphix-dct-data-bookmarks:x.0.0   "/usr/bin/tini -- ./..."           7 seconds
ago          Up 5 seconds        delphix-dct-data-bookmarks:x.0.0
20d1782cb3bb  delphix-dct-ui:x.0.0               "node ./index.js"                   7 seconds
ago          Up 5 seconds        delphix-dct-ui:x.0.0
4fae43c79e8d  delphix-dct-virtualization:x.0.0   "/usr/bin/tini -- ./..."           7 seconds
ago          Up 5 seconds        delphix-dct-virtualization:x.0.0
83d7d661d8a0  delphix-dct-graphql:x.0.0         "/bin/sh -c 'BASE_UR..."           7 seconds
ago          Up 6 seconds        delphix-dct-graphql:x.0.0
3dded474e28b  delphix-dct-postgres:x.0.0        "docker-entrypoint.s..."           7 seconds
ago          Up 6 seconds        5432/tcp delphix-dct-postgres:x.0.0
2sdr89y896h   delphix-dct-masking:x.0.0         "java -jar /opt/delp..."           7 seconds
ago          Up 5 seconds
```

⁴⁷ <https://download.delphix.com/folder>

6.3.1.5 Initial logging configuration


To restrict boundless log file generation, DCT is now shipped with default logging drivers provided with default values for `max-size` and `max-file` options, for all services in the `docker-compose.yaml` file. These values can also be overwritten to correspond with the user requirement. Use the process below to modify the settings.

1. Open the `docker-compose.yaml` file.
2. Add the following lines in every service section below.

```
logging:
  driver: "json-file"
  options:
    max-file: "5"
    max-size: 10m
```

- a. The `driver` parameter above means the logging driver is a JSON file. For these services, a maximum of five log files (`max-file`) will be retained with a size of 10MB (`max-size`), after which files are overwritten.
- b. The maximum number of logs and the maximum log file size shown in 2a can be changed by updating the values for `max-file` and `max-size`.

6.3.2 Bootstrapping API Keys

 Delphix has announced the deprecation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month deprecation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.


It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁴⁸ or [OpenShift](#)⁴⁹.

API keys are the default method to authenticate with DCT. This is done by including the key in the [HTTP Authorization request header](#)⁵⁰ with type `apk`.

⁴⁸ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁴⁹ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

⁵⁰ <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Authorization>


 API keys are long-live tokens and as a result, do not automatically expire in the future. They remain valid until they are deleted or destroyed from DCT.

A cURL example using an example key of

`1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3` would appear as:

```
curl --header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3'
```

cURL (like web browsers and other HTTP clients) will not connect to DCT over HTTPS unless a valid TLS certificate has been configured for the Nginx server. If this [configuration step](#) (see page 116) has not been performed yet and the risk is comprehended, you may disable the check in the HTTP client. For instance, this can be done with cURL using the `--insecure` flag.

 The cURL version must be 7.43 or higher.

6.3.2.1 Bootstrap first API key

There is a special process to bootstrap the creation of the first API key. This first API key should only be used to create another key and then promptly deleted, since the bootstrap API will appear in the logs. This process can be repeated as many times as needed, for example, in a case where existing API keys are lost or have been deleted. It also means that the Linux users with permissions to edit the docker-compose file implicitly have the ability to get an API key at any time. There is no mechanism to lock this down after the first bootstrap key is created.

Begin by stopping the application with the following command:

```
sudo docker-compose stop
```

Once the application is stopped, edit the `docker-compose.yml` file and modify the following lines to the DCT section, to set the `API_KEY_CREATE` to the string value "true":

```
services:
  gateway:
    environment:
      API_KEY_CREATE: "true"
```

Start DCT again with `sudo docker-compose up`. You will see the following output in the logs for the app container (the key will be different from this example):


```
NEWLY GENERATED API KEY: 1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3
```

Copy the API Key and shut down the DCT app. The API key can now be used to authenticate with DCT. Remember that the API Key value must be prefixed with apk. An example cURL command with the above API Key appears as follows:

```
curl --header 'Authorization: apk  
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3'
```

Edit the `docker-compose.yaml` file to set the `API_KEY_CREATE` environment variable value back to "false" and restart DCT again with `sudo docker-compose up -d`.

6.3.3 Custom configuration



Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁵¹ or [OpenShift](#)⁵².

6.3.3.1 Introduction

DCT was designed for users to configure Delphix applications in a way that would meet their security requirements, which handled with a custom configuration. This article provides background information on performing custom configurations, which are referenced throughout DCT articles and sections.

6.3.3.2 Bind mounts

Configuration of DCT is achieved through a combination of API calls and the use of Docker [bind mounts](#)⁵³. A bind mount is a directory or file on the host machine that will be mounted inside the container. Changes made to the files on the host machine will be reflected inside the container. It does not matter where the files live on the host machine, but the files must be mounted to specific locations inside the container so that the application can find them.

⁵¹ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁵² <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

⁵³ <https://docs.docker.com/storage/bind-mounts/>

The DCT and proxy containers can both be configured via separate bind mounted directories. Each container requires all configuration files to be mounted to the `/etc/config` directory inside the container. Therefore, it is recommended to create a directory for each container on the host machine to store all of the configuration files and mount them to `/etc/config`. This is done by editing the `docker-compose.yml`. Under **proxy services**, add a **volumes** section if one does not already exist; this is used to mount the configuration directory on the host to `/etc/config`. For example, if `/my/proxy/config` is the directory on the host that contains the configuration files, then the relevant part of the compose file would look like this:

```
services:
  proxy:
    volumes:
      - /my/proxy/config:/etc/config
```

To change the configuration of the DCT container, make a similar change under its service section, the only difference being the directory on the host. After making this change, the application will need to be stopped and restarted.

The structure of `/my/proxy/config` will need to match the required layout in `/etc/config`. When each container starts, it will create default versions of each file and place them in the expected location. It is highly recommended to start from the default version of these files. For example, if `/my/proxy/config` is the bind mount directory on the host, it could be populated with all the default configuration files by running the following commands.

First, create an `nginx` directory inside `/my/proxy/config` on the host.

```
cd /my/proxy/config
mkdir nginx
```

Find the **id** of the proxy container with **docker ps**. Look for the container with a **delphix-dct-proxy** image name. To determine the user and group ownership for any configuration files, start the containers and open a shell to the relevant one (`nginx` in this example), then examine the current user/group IDs associated with the files (where `x.0.0` should be changed to the version of DCT being installed).

```
docker ps
```


CONTAINER ID	IMAGE	COMMAND	CREATED
ac343412492a	delphix-dct-proxy:x.0.0	"/bootstrap.sh"	8 minutes ago
8 minutes	0.0.0.0:443->443/tcp, :::443->443/tcp	dct-packaged_proxy_1	Up

In the above example, `ac343412492a` is the **id**. Run the following command to copy the default files to the bind mount.

```
docker cp <container id>:/etc/config/nginx /my/proxy/config/nginx
```

One can always go back to the original configuration by removing the bind-mount and restarting the container or using `docker cp` as in the previous example to overwrite the custom files with the default versions.

6.3.4 Docker logs

 Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁵⁴ or [OpenShift](#)⁵⁵.

DCT leverages the [Docker logging](#)⁵⁶ infrastructure. All containers log to `stdout` and `stderr` so that they are processed by Docker. Docker supports logging drivers for a variety of tools, such as Fluentd, Amazon CloudWatch, and Splunk.

Use the [Configure logging drivers](#)⁵⁷ Docker article to configure them – these changes will need to be made to the `docker-compose.yml` file. This [section of the Docker article](#)⁵⁸ explains how to alter the Compose file to adjust the logging driver. For example, to use `syslog` for the proxy container, it would appear as follows.

```
services:
  proxy:
    logging:
      driver: syslog
      options:
        syslog-address: "tcp://192.123.1.23:123"
```

6.3.5 Migration topics

- [Migrate to Kubernetes](#) (see page 84)
- [Migrate to OpenShift](#) (see page 87)

54 <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

55 <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

56 <https://docs.docker.com/config/containers/logging/>


57 <https://docs.docker.com/config/containers/logging/configure/>

58 <https://docs.docker.com/compose/compose-file/compose-file-v3/#logging>

6.3.5.1 Migrate to Kubernetes

6.3.5.1.1 Overview

Installations starting on Docker Compose may be migrated to Kubernetes by moving the persistent data store using the following steps. In-place upgrades from Docker Compose to Kubernetes are not supported.

 During the migration process, there will be a downtime period where the service cannot be used.

6.3.5.1.1.1 Migration Process

Stop DCT services. In order to avoid a situation of losing data, stop serving the upcoming traffic with:

```
docker-compose stop
```

Copy the Postgres Docker volume folder data on a local machine with:


```
mkdir database
docker cp {dbcontainer_Id}:/var/lib/postgresql/data ./database
```

Copy the encryption key Docker volume folder data on a local machine with:

```
mkdir data_key
docker cp {gateway_container_id}:/data ./data_key
```

Copy the masking encryption key Docker volume folder data on a local machine with:

```
mkdir masking_data_key
docker cp {masking_container_id}:/data ./masking_data_key
```

 Mounted Docker volume folder content for database is copied in database folder on local machine.

Mounted Docker volume folder content for encryption key is copied in the data_key folder on local machine.

Mounted masking Docker volume folder content for encryption key is copied in the `masking_data_key` folder on local machine.

Move the copied volume folders (**database, data_key and masking_data_key** from the previous step) to the Kubernetes host machine where DCT is up and running.

Update the **values.yaml** file to add the list of certificates which were used in the previous DCT version (present in mounted trustStore). Update the deployment with the new **values.yaml** file.

Terminate the proxy pod to stop serving external traffic with:

```
kubectl scale --replicas=0 deployment/proxy -n dct-services
```

Terminate the database to stop internal threads using the database with:

```
kubectl scale --replicas=0 deployment/database -n dct-services
```

Create a dummy pod to access the Persistent Volume. Use the Pod.yaml as an example:

```
apiVersion: v1
kind: Pod
metadata:
  Namespace: dct-services
  name: dummy-pod
  labels:
    app: dummy-pod
spec:
  containers:
    - image: ubuntu
      command:
        - "sleep"
        - "604800"
      imagePullPolicy: IfNotPresent
      name: ubuntu
  restartPolicy: Always
  volumes:
    - name: gwdatabase-data
      persistentVolumeClaim:
        claimName: gwdatabase-data
```

Followed by this command to actually create the dummy pod:

```
kubectl apply -f pod.yaml -n dct-services
```

Restore previous DCT version volume data with DCT deployed on the Kubernetes setup (in Persistent Volume).

Move the encryption key with:

```
cd data_key
kubectl cp data dct-services/{gateway_pod_name}:/
```

Move the masking encryption key with:

```
cd masking_data_key
kubectl cp data dct-services/{masking_pod_name}:/
```

Move the Postgres data with:

```
cd database
kubectl cp data dct-services/{dummy_pod_name}:/var/lib/postgresql
```

Delete the dummy pod with:

```
kubectl delete pod dummy-pod -n dct-services
```

Start the database pod (scale to 1) with:

```
kubectl scale --replicas=1 deployment/database -n dct-services
```

Delete or patch the gateway pod with:

```
kubectl delete pod {gateway_pod_name} -n dct-services
```

Delete or patch the data-library pod with:

```
kubectl delete pod {data-library_pod_name} -n dct-services
```

Delete or patch the jobs pod with:

```
kubectl delete pod {jobs_pod_name} -n dct-services
```

Delete or patch the data-bookmarks pod with:

```
kubectl delete pod {data-bookmarks_pod_name} -n dct-services
```

Delete or patch the masking pod with:

```
kubectl delete pod {masking_pod_name} -n dct-services
```

Start the proxy service to serve the external service:

```
kubectl scale --replicas=1 deployment/proxy -n dct-services
```

6.3.5.2 Migrate to OpenShift

6.3.5.2.1 Overview

Installations starting on Docker Compose may be migrated to OpenShift by moving the persistent data store using the following steps. In-place upgrades from Docker Compose to OpenShift are not supported.



During the migration process, there will be a downtime period where the service cannot be used.

6.3.5.2.1.1 Migration Process

Stop DCT services. In order to avoid a situation of losing data, stop serving the upcoming traffic with:

```
docker-compose stop
```

Copy the Postgres Docker volume folder data on a local machine with:

```
mkdir database  
docker cp {dbcontainer_Id}:/var/lib/postgresql/data ./database
```

Copy the encryption key Docker volume folder data on a local machine with:

```
mkdir data_key  
docker cp {gateway_container_id}:/data ./data_key
```

Copy the masking encryption key Docker volume folder data on a local machine with:

```
mkdir masking_data_key  
docker cp {masking_container_id}:/data ./masking_data_key
```

- Mounted Docker volume folder content for database is copied in database folder on local machine.
- Mounted Docker volume folder content for encryption key is copied in the `data_key` folder on local machine.
- Mounted masking Docker volume folder content for encryption key is copied in the `masking_data_key` folder on local machine.

Move the copied volume folders (**database, data_key and masking_data_key** from the previous step) to the Kubernetes host machine where DCT is up and running.

Update the **values.yaml** file to add the list of certificates which were used in the previous DCT version (present in mounted trustStore). Update the deployment with the new **values.yaml** file.

Terminate the proxy pod to stop serving external traffic with:

```
oc scale --replicas=0 deployment/proxy -n dct-services
```

Terminate the database to stop internal threads using the database with:

```
oc scale --replicas=0 deployment/database -n dct-services
```

Create a dummy pod to access the Persistent Volume. Use the Pod.yaml as an example:

```
apiVersion: v1
kind: Pod
metadata:
  Namespace: dct-services
  name: dummy-pod
  labels:
    app: dummy-pod
spec:
  containers:
    - image: ubuntu
      command:
        - "sleep"
        - "604800"
      imagePullPolicy: IfNotPresent
      name: ubuntu
  restartPolicy: Always
  volumes:
    - name: gwdatabase-data
      persistentVolumeClaim:
        claimName: gwdatabase-data
```

Followed by this command to actually create the dummy pod:


```
oc apply -f pod.yaml -n dct-services
```

Restore previous DCT version volume data with DCT deployed on the Kubernetes setup (in Persistent Volume).

Move the encryption key with:

```
cd data_key
oc cp data dct-services/{gateway_pod_name}:/
```

Move the encryption key with:

```
cd masking_data_key
oc cp data dct-services/{masking_pod_name}:/
```

Move the Postgres data with:

```
cd database
oc cp data dct-services/{dummy_pod_name}:/var/lib/postgresql
```

Delete the dummy pod with:

```
oc delete pod dummy-pod -n dct-services
```

Start the database pod (scale to 1) with:

```
oc scale --replicas=1 deployment/database -n dct-services
```

Delete or patch the gateway pod with:

```
oc delete pod {gateway_pod_name} -n dct-services
```

Delete or patch the data-library pod with:

```
oc delete pod {data-library_pod_name} -n dct-services
```

Delete or patch the jobs pod with:

```
oc delete pod {jobs_pod_name} -n dct-services
```

Delete or patch the data-bookmarks pod with:

```
oc delete pod {data-bookmarks_pod_name} -n dct-services
```

Delete or patch the masking pod with:

```
oc delete pod {masking_pod_name} -n dct-services
```

Start the proxy service to serve the external service:

```
oc scale --replicas=1 deployment/proxy -n dct-services
```

6.3.6 Admin topics for Docker Compose

- [Backup DCT on Docker Compose \(see page 90\)](#)
- [Deployment upgrade for Docker Compose \(see page 91\)](#)
- [Factory reset DCT for Docker Compose \(see page 93\)](#)

6.3.6.1 Backup DCT on Docker Compose



Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁵⁹ or [OpenShift](#)⁶⁰.

This article discusses how to backup DCT. The data that needs to be backed up is the Docker volumes used by the DCT container, gwdatabase container, and the configuration directories on the host that are bind mounted to the containers.

The Docker volumes named `{xxx}delphix-dct-data` , `{xxx}delphix-dct-database-data` and `{xxx}delphix-dct-masking-data` should be backed up to prevent data loss. This [Docker article](#)⁶¹ explains how to backup a data volume.

The bind mount directories containing the configuration files are standard directories that can be backed up as desired. A simple approach would be to create a tar file of the contents. If `/my/config` is the bind mount directory on the host, then this can be done with the following command:


⁵⁹ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁶⁰ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>


⁶¹ <https://docs.docker.com/storage/volumes/#backup-restore-or-migrate-data-volumes>

```
tar -czf gateway-backup.tgz /my/config
```

6.3.6.2 Deployment upgrade for Docker Compose

 Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.


It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁶² or [OpenShift](#)⁶³.

 DCT versions 2.0.0 through 6.0.2 running on Docker Compose, that are being upgraded to DCT 7.0.0 or later, may experience potential failure to start post-upgrade, resulting in a "permission denied" error in the logs. Operations post-upgrade may also fail with internal errors. The issue is due to the UID running the application containers changing from UID 1000 (in DCT 2.0.0 through 6.0.2) to UID 1010 (in DCT 7.0.0 and later). Resolving the issues requires the following one-time change and no container restart is required:

1. Change ownership of the volume associated to the gateway container to the new UID:


```
docker exec -u 0 -it <gateway-container-name> chown delphix:delphix /data
```
2. If bind mounts have been used to configure DCT, they must grant permission to the user with UID 1010 (GUID 1010) to read/write files, for example:


```
chown 1010:1010 /path/to/nginx/bind/mount
```

 Starting from DCT 12.0.0 we have introduced one new service (container) named masking. Memory and CPU requirement for this new container is 512Mi and 500m respectively.

This container creates a new volume mount `masking-data`.

If you are using masking engines with self sign certificated and previously mounted the truststore to the gateway pod like:

⁶² <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁶³ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

```
volumes:
- gateway-data:/data
- /truststore/config/path/on/host:/etc/config/certs
```

Then please mount this old truststore to the masking service as well by keeping it in separate folder and uncomment below lined in docker-compose.yaml masking service section:

```
volumes:
- masking-data:/data
- /truststore/config/path/on/host:/etc/config/certs
```

6.3.6.2.1 Introduction

This article describes the procedure to upgrade the DCT version without losing any data. Docker Compose uses the concept of 'project' to create unique identifiers for all of a project's containers and other resources (like volumes, etc.).

Get the current project name and note it down using the following command:

The volume name would be of the format **{project-name}_gateway-data** and **{project-name}_gwdatabase-data**. In the below example, the project name is **delphix-dct**.

```
docker volume ls
DRIVER      VOLUME NAME
local      delphix-dct_gateway-data
local      delphix-dct_gwdatabase-data
local      delphix-dct_masking-data
```

Bring down DCT services using the following command:

```
docker compose down
```

Refer to the **Installation and Setup** article to download and extract the new release tarball, then load Docker images.

Navigate to the extracted directory which contains the **docker-compose.yaml** file. By default, Docker Compose uses the extracted folder name as **project-name**.



Edit the docker-compose.yaml file. Changes made to the docker-compose.yaml prior to upgrade file must be applied to the newly extracted docker-compose.yaml file.

With that, either rename the extracted folder to match the project-name and run:

```
docker compose up -d
```

OR run the below command with the project-name noted above from step #1 above

```
docker compose -p <project-name> up -d
```



If the `-p` argument is used to deploy DCT services, then the corresponding command to bring down the DCT services would be:

```
docker compose -p <project-name> down
```

6.3.6.3 Factory reset DCT for Docker Compose



Delphix has announced the depreciation of support for Docker Compose with DCT version 13.0.0. The January 2024 release starts a 12-month depreciation period for all supported versions on Docker Compose. All prior and current product versions will continue to be supported on Docker Compose until January 2025. However, production DCT workloads in Docker Compose are not supported.

It is highly recommended that new DCT installations are performed on [Kubernetes](#)⁶⁴ or [OpenShift](#)⁶⁵.

This article explains how to factory reset DCT. Factory resetting means deleting all of the configuration and data associated with DCT. Perform this step only if you are absolutely sure about this and understand the implications.

Bring all of the DCT services down with this command:

```
docker compose down
```

List all Docker volumes being used and note down the volume names:

```
docker volume ls
DRIVER  VOLUME NAME
local   dct_gateway-data
local   dct_masking-data
```

⁶⁴ <https://dct.delphix.com/docs/latest/installation-and-setup-for-kubernetes>

⁶⁵ <https://dct.delphix.com/docs/latest/installation-and-setup-for-openshift>

```
local    dct_gwdatabase-data
```

Delete the Docker volumes that are listed from the previous command:

```
docker volume rm dct_gateway-data
docker volume rm dct_gwdatabase-data
docker volume rm dct_masking-data
```

6.4 Engines: connecting/authenticating

6.4.1 Introduction

After DCT Authentication is complete, the HTTPS should be securely configured on DCT and able to be authenticated against. The next step is to register an engine with DCT so that it can fetch results. DCT connects to all engines over HTTPS, thus some configurations might be required to ensure it can communicate successfully.

6.4.2 Truststore for HTTPS

If the CA certificate that signed the engine's HTTPS certificate is not a trusted root CA certificate present in the JDK, then custom CA certificates can be provided to DCT. If these certificates are not provided, a secure HTTPS connection cannot be established and registering the engine will fail. The `insecure_ssl` engine registration parameter can be used to bypass the check, however, this should not be used unless the risks are understood.

Get the public certificate of the CA that signed the engine's HTTPS certificate in PEM format. IT team help may be required to get the correct certificates. Base64 encode the certificate with:

```
cat mycertfile.pem | base64 -w 0
```

Copy the Base64 encoded value from the previous step and configure in `values.yaml` file under `truststoreCertificates` section. e.g. section will look like this:

```
truststoreCertificates:
  <certificate_name>.cert: <base64 encode certificate string value in single line>
```

<certificate_name> can be any logically valid string value for e.g. "engine".

All the certificates configured in `truststoreCertificates` section will be read and included in the `trustStore` which would be then used for SSL/TLS communication between DCT and Delphix Engine.

6.4.3 Authentication with engine

All authentication with the Delphix Engine is done with the username and password of a domain admin engine user. There are two methods of storing these credentials with DCT. They can either be stored and encrypted on DCT itself or retrieved from a password vault. We recommend fetching the credentials from a vault. Currently only the HashiCorp vault is supported.

6.4.4 HashiCorp vault

There are two high-level steps to configuring a HashiCorp vault. The first is to set up authentication with the vault and register the vault. The second is to tell DCT how to get the specific engine credentials needed from that registered vault. A single vault can be used for multiple different Delphix Engines.

6.4.4.1 Vault authentication and registration

First, DCT needs to be able to authenticate with the vault. DCT supports the [Token](#)⁶⁶, [AppRole](#)⁶⁷, and [TLS Certificates](#)⁶⁸ authentication methods. This is done by passing a command to the [HashiCorp CLI](#)⁶⁹. It is recommended to first ensure that successful authentication is done and one can retrieve the credentials with the HashiCorp CLI directly to ensure the correct commands are passed to DCT.

Adding a vault to DCT is done through API calls to the `/v2/management/vaults/hashicorp` endpoint. All authentication methods requires the location of the vault is provided through the `env_variables` property in the POST body like so:

```
"env_variables": {
  "VAULT_ADDR": "https://10.119.132.40:8200"
}
```

6.4.4.2 Token

To use the token authentication method, this needs to be included as part of the `env_variables` field. The full example to register the vault would appear as:

```
curl --location --request POST 'https://<hostname>/v2/management/vaults/hashicorp' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "env_variables": {
    "VAULT_TOKEN": "<your token>"
    "VAULT_ADDR": "https://10.119.132.40:8200"
  }
}
```

66 <https://www.vaultproject.io/docs/auth/token>

67 <https://www.vaultproject.io/docs/auth/approle>

68 <https://www.vaultproject.io/docs/auth/cert>

69 <https://www.vaultproject.io/docs/commands>

```
}
}'
```

A response should be received similar to the lines below:

```
{
  "id": 2,
  "env_variables": {
    "VAULT_TOKEN": "<your token>"
    "VAULT_ADDR": "https://10.119.132.40:8200"
  }
}
```

Note the id of the vault, this will be needed in the next step to register the engine.

6.4.4.3 AppRole

To use the AppRole authentication method, this needs to be included as part the login_command_args field, as shown below.

```
"login_command_args":
  [ "write", "auth/approle/login", "role_id=1", "secret_id=123"]
```

The full example to register the vault would appear as:

```
curl --location --request POST 'https://<hostname>/v2/management/vaults/hashicorp' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "env_variables": {
    "VAULT_ADDR": "https://10.119.132.40:8200"
  },
  "login_command_args":
    [ "write", "auth/approle/login", "role_id=1", "secret_id=123"]
}'
```

A response should be received similar to the lines below:

```
{
  "id": 2,
  "env_variables": {
    "VAULT_TOKEN": "<your token>"
    "VAULT_ADDR": "https://10.119.132.40:8200"
  }
}
```


6.4.5 TLS certificates

The configuration of mutual TLS authentication requires an additional step. This feature currently is NOT supported for Kubernetes deployment of DCT. This will be covered in later releases.

6.4.5.1 Retrieving engine credentials

Once DCT can authenticate with the vault, it needs to know how to fetch the relevant engine credentials. When registering an engine, the user will need to provide the HashiCorp CLI commands through the `hashicorp_vault_username_command_args` and `hashicorp_vault_password_command_args` parameters.


The relevant part of the engine registration payload will look like the following:

```
'{
  "hashicorp_vault_id": 1
  "hashicorp_vault_username_command_args": ["kv", "get", "--field=username", "kv-
v2/delphix-engine-secrets/engineUser"]
,
  "hashicorp_vault_password_command_args": ["kv", "get", "--field=password", "kv-
v2/delphix-engine-secrets/engineUser"]
}'
```

The `hashicorp_vault_id` will be the ID that was returned as part of the previous step. Note that the exact paths to fetch the username and password will vary depending on the exact configuration of the vault.

6.5 Accounts: connecting/authenticating

There are 5 supported methods for authentication; **API keys**, **Username/Password**, **LDAP/Active Directory**, **SAML/SSO**, and **OpenID Connect**. These authentication methods are detailed on the corresponding pages in this section.

 DCT uses Nginx/[OpenResty](https://openresty.org/en/)⁷⁰ as an HTTP server and a reverse proxy for the application. Using the default configuration, all connections to DCT are over HTTPS and require the user to authenticate. The Nginx/OpenResty configuration files can be edited via `/etc/config` bind mounts, for the proxy container to customize the HTTP server and change options (such as TLS versions).

- [API keys](#) (see page 98)
- [Username/password](#) (see page 100)
- [LDAP/Active Directory](#) (see page 103)

⁷⁰ <https://openresty.org/en/>

- [SAML/SSO \(see page 111\)](#)

6.5.1 API keys

API keys are the default method to authenticate with DCT. This is done by including the key in the [HTTP Authorization request header](#)⁷¹ with type `apk`.



API keys are long-live tokens and as a result, do not automatically expire in the future. They remain valid until they are deleted or destroyed from DCT.

A cURL example using an example key of

`1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3` would appear as:

```
curl --header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3'
```

cURL (like web browsers and other HTTP clients) will not connect to DCT over HTTPS unless a valid TLS certificate has been configured for the Nginx server. If this [configuration step \(see page 116\)](#) has not been performed yet and the risk is comprehended, you may disable the check in the HTTP client. For instance, this can be done with cURL using the `--insecure` flag.




The cURL version must be 7.43 or higher.

6.5.1.1 Create and manage API Keys

The initial API key created should be used to create a new admin secure key. This is done by creating a new Account entity and setting the `generate_api_key`. The "username" attribute should be the desired name to uniquely identify the account.

⁷¹ <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Authorization>

```
curl --location --request POST 'https://<hostname>/v2/management/accounts' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3' \
--data-raw '{
  "username": "secure-key",
  "generate_api_key": true
}'
```

 If the cURL version being used is below 7.43, replace the `--data-raw` option with `--data` .

A response should be received similar to the lines below:

```
{
  "id": 2,
  "token": "2.vCfC0MnpySYZLshuxap2aZ7xqBKAnQvV7hFnobe7xuNlHS9AF2NqnV9XXw4UyET6"
  "username": "secure-key"
}
```

Now that the new and secure API key is created, the old one must be deleted for security reasons since the key appeared in the logs. To do this make the following request:

```
curl --location --request DELETE 'https://<hostname>/v2/management/api-clients/<id>' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk
2.vCfC0MnpySYZLshuxap2aZ7xqBKAnQvV7hFnobe7xuNlHS9AF2NqnV9XXw4UyET6'
```

The id referenced above is the numeric id of the Account. It is the integer before the period in the token. For example, the id of

`1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3` is 1.


Finally, to list all of the current Accounts, make the following request:

```
curl --location --request GET 'https://<hostname>/v2/management/accounts/' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>'
```

6.5.2 Username/password

When creating an account, a username and password combination can be associated with the account (whether an API Key was generated for the account or not). To do so, specify the “username” and “password” properties in the API request, for example:

```
curl -k --location --request POST 'https://<hostname>/v2/management/accounts' \
  --header 'Content-Type: application/json' \
  --header 'Accept: application/json' \
  --header 'Authorization: apk
1.0p9PMkZ04Hgy0ezwjhX0Fi4lEKrD4pflejgqjd0pfKtywLSWR9G0fIaWajuKcBT3' \
  --data-raw '{
  "username": "some-username",
  "password": "some-password",
  "generate_api_key": false
  "is_admin": true
}'
```

 The **is_admin** property will create the account with admin privileges. Remove this property to create an account without admin privileges.

The username and password combination can then be used to login via the UI, or to fetch a temporary access token valid for 24 hours. To do so, call the ‘login’ API endpoint:

```
curl -k --location --request POST 'https://<hostname>/v2/login' \
  --header 'Content-Type: application/json' \
  --header 'Accept: application/json' \
  --data-raw '{
  "username": "some-username",
  "password": "some-password"
}'
```

A response should be received similar to the lines below:

```
{
  "access_token": "eyJhbGciOiJIUzI1NiJ9.eyJpc3MiOiJhcGlndy1zZXJ2aWNlcylhcHAiLCJzdWIiOiI4IiwiaWF0IjoiYyNTUyMzI3LCJpYXQiOiJlE2NjI0NjU5MjcsInVzZXJlIjoic29tZS1lc2VybmFtZSJ9.Cx_hGU9noyWS6mtK6gjsA85FTgJRQgyJizR5t_akNps",
  "token_type": "Bearer",
  "expires_in": 86400
}
```

The access token can be used as [HTTP Authorization request header](#)⁷² with type **Bearer**.

i This access token should not be confused with ApiKeys. These tokens are short lived tokens and expire at the end of 24 hours.

A cURL example using the access token retrieved above would appear as:

```
curl --header 'Authorization: Bearer
eyJhbGciOiJIUzI1NiJ9.eyJpc3MiOiJhcGlnZy1zZXJ2aWNlcy1hcHAiLCJzdWIiOiI4IiwiaXhwIjoxNjYy
NTUyMzI3LCJpYXQiOjE2NjI0NjU5MjcsInVzZXJlIjoic29tZS11c2VybmFtZSJ9.Cx_hGU9noyWS6mtK
6gjsA85FTgJRQgyJizR5t_akNps'
```

The password for an account can be updated with the **change_password** API endpoint, passing in both the old and new passwords, such as in this example:

```
curl -k --location --request POST '<hostname>/v2/management/accounts/3/
change_password \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: Bearer
eyJhbGciOiJIUzI1NiJ9.eyJpc3MiOiJhcGlnZy1zZXJ2aWNlcy1hcHAiLCJzdWIiOiI4IiwiaXhwIjoxNjYy
NTUyMzI3LCJpYXQiOjE2NjI0NjU5MjcsInVzZXJlIjoic29tZS11c2VybmFtZSJ9.Cx_hGU9noyWS6mtK
6gjsA85FTgJRQgyJizR5t_akNps' \
--data-raw '{
  "old_password": "some-password",
  "new_password": "new-password"
}'
```

Following security best practices, the password is not stored on DCT and cannot be retrieved. If the password has been lost, an account with admin privilege can reset the password for a particular account. It is recommended to change the password reset by an admin account on the first login, or with the **change_password** API, as described above.

⁷² <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Authorization>

```
curl -k --location --request POST '<hostname>/v2/management/accounts/2/
password_reset' \
  --header 'Content-Type: application/json' \
  --header 'Accept: application/json' \
  --header 'Authorization: Bearer
eyJhbGciOiJIUzI1NiJ9.eyJpc3MiOiJhcGlndy1zZXJ2aWNlcy1hcHAiLCJzdWIiOiI4IiwiaXhwIjoxNjYy
NTUyMzI3LCJpYXQiOjE2NjI0NjU5MjcsInVzZXJ1Ym91Ijoic29tZS1lc2VybmFtZSJ9.Cx_hGU9noyWS6mtK
6gjsA85FTgJRQgyJizR5t_akNps' \
  --data-raw '{
    "new_password": "new-password"
  }'
```

In the above example, the admin is resetting the password of an account with id **2** to “new-password”.

6.5.2.1 Password policies

The password policy feature allows users to enable and customize the password policy enforced for local username/password authentication (does not apply to LDAP/Active Directory or SAML/SSO based authentication).

6.5.2.2 Understanding password policies

The password policy is a set of requirements that local passwords must satisfy.

- **min_length**: A password must be longer than this length.
- **reuse_disallow_limit**: The user should not reuse old passwords. This tells the number of last used passwords disallowed to be reused as the new passwords.
- **uppercase_letter**: A password must have at least one capital letter.
- **lowercase_letter**: A password must have at least one lower case letter.
- **digit**: A password must have at least one digit.
- **special_character**: A password must have at least one special character, such as #, \$, !
- **disallow_username_as_password**: A password should not be the same as the user name.
- **maximum_password_attempts**: The number of allowed attempts for incorrect password, after which the account gets locked.

6.5.2.3 Default password policy

By default, DCT does not enforce any password policy.

6.5.2.4 Changing the password policy

To change the current password policy, call the password policy API endpoint, as shown in the example below:

```
curl --location --request PATCH 'https://<hostname>/v2/management/accounts/password-
policies' \
```

```

--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "enabled": true,
  "maximum_password_attempts": 2,
  "min_length": 5,
  "reuse_disallow_limit": 3,
  "digit": true,
  "uppercase_letter": true,
  "lowercase_letter": true,
  "special_character": true,
  "disallow_username_as_password": true
}'

```

Changing the password policy does not affect existing passwords.

6.5.2.5 Disabling local username/password authentication

Username/password authentication (with passwords locally in DCT) can be disabled for individual accounts by not setting or unsetting their password property, or across the DCT instance using the global properties API. Disable username/password authentication to force authentication to use an alternate authentication method (LDAP/Active Directory, SAML/SSO, etc.) as shown in this example:

```

curl --location --request PATCH 'https://<hostname>/v2/management/properties' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{"disable_username_password": true}'

```

6.5.3 LDAP/Active Directory

6.5.3.1 Configuration

LDAP/Active directory can be used to authenticate login requests, and optionally to retrieve additional information about accounts, thereafter referred to as LDAP Search.

6.5.3.1.1 Configuring authentication

The following attributes must be set to configure LDAP/Active Directory authentication.

Property Name	Description
enabled	Whether the LDAP/Active Directory feature is enabled.

Property Name	Description
auto_create_users	<p>Whether DCT must automatically create account records for successful authentication attempts using a username which does not match any accounts.</p> <p>If this is disabled, an administrator must create a DCT account with an ldap_principal attribute matching the value from the LDAP/Active Directory server prior to the first login attempt.</p> <p>If this is enabled, any user with valid credentials in the LDAP/Active Directory server can authenticate to DCT, by default with an empty authorization set (i.e not being able to view any data or perform any action).</p>
hostname	The host name or IP address of the LDAP/Active Directory server.
port	Port of the LDAP/Active Directory server. This is usually 389 for non SSL, and 636 for SSL.
enable_ssl	Whether the connection to the LDAP/Active Directory server must be performed over SSL. It is highly advised to use SSL. Without SSL, communication between DCT and the LDAP/Active server can be intercepted.
insecure_ssl, unsafe_ssl_hostname_check, trustore_file_name, truststore_password	The SSL protocol requires the LDAP/Active Directory server to expose a certificate signed by a Certificate Authority (CA) trusted by the JDK which is running DCT. Refer to the dedicated section below to see how to configure an Active Directory/ LDAP server of which certificate is not recognized.
[domains].msad_domain_name	<p>Microsoft Active Directory only: The DNS name of a domain in the same forest as the accounts which login. DCT will append the msad_domain_name to the username provided at login to form a user principal name (UPN).</p> <p>Example: if the msad_domain_name is http://mycompany.co and a user logs in with username john, DCT will perform an LDAP request to the</p>

Property Name	Description
	Active Directory server to authenticate john@mycompany.co ⁷³ .
[domains].username_pattern	<p>If the LDAP server is not Microsoft Active Directory, the username_pattern is used to create a DN string for user authentication. The pattern argument {0} is replaced with the username at runtime.</p> <p>Example: If the username_pattern is uid={0},ou=People and a user logs in with username john, DCT will perform an LDAP request with DN uid=john,ou=People.</p>

The LDAP/Active Directory Integration can be configured both via DCT UI and API. The below image shows an example of how the configuration can be set in the UI as a way to Authenticate users, auto create new users, as well as map group attributes for authorization within the DCT Access Control system.

⁷³ <mailto:john@mycompany.co>

Edit LDAP Settings

Enabled

Auto-create Users

Hostname
activedirectory.acme.com

Port
636

Domains +

MSAD Domain Name
acme.com

Username Pattern

Search Base
CN=Users,DC=acme,DC=com

Group Attribute
department

Email Attribute
mail + 🗑️

First Name Attribute
givenName

Last Name Attribute
sn

Object Class Attribute
person

Search Attribute
sAMAccountName

Cancel Save

The following example requests enable LDAP authentication over SSL with an Active Directory server at address **activedirectory.company.co** using the **us.company.co** domain:

```
curl --location --request PUT 'https://<hostname>/v2/management/ldap-config' \
```

```

--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "enabled": true,
  "auto_create_users": true,
  "hostname": "activedirectory.company.co",
  "enable_ssl": true,
  "port": 636,
  "domains": [{
    "msad_domain_name": "us.company.co"
  }]
}'

```

6.5.3.1.2 Validating the configuration

Updating the LDAP/Active Directory configuration does not guarantee that the provided values are correct, as validating those requires a user to authenticate to DCT. This can be achieved with the `ldap-config/validate` API endpoints, using the credentials valid for the LDAP/Active Directory server. When provided with a username/password combination, the `ldap-config/validate` API endpoint will authenticate with the LDAP server. If the response status code is 200, the configuration is correct. Otherwise, the response code will be 400, and the response body will provide information to resolve the configuration problems. For example:

```

curl --location --request POST 'https://<hostname>/v2/management/ldap-config/validate' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "username": "<ldap-username>",
  "password": "<ldap-password>"
}'

```



Because of a defect in version 3.0.0 of DCT, the above request might fail with a response similar to:

```

search failed for john.doe with search base null' ,search attribute
'null'

```

This indicates that authentication works, and search (see below) is not configured.

6.5.3.1.3 Login

Once the configuration has been updated, accounts can login (via the UI or API) using the same UI form/API endpoint they would be using for the local username/password authentication feature. For example:

```
curl -k --location --request POST 'https://<hostname>/v2/login' \  
--header 'Content-Type: application/json' \  
--header 'Accept: application/json' \  
--data-raw '{  
  "username": "<ldap-username>",  
  "password": "<ldap-password>"  
}'
```

When LDAP/Active directory is enabled, DCT first attempts to validate passwords with the LDAP/Active Directory server, and falls back to local password authentication in case of failure. Enabling LDAP/Active directory is thus a non disruptive operation for existing accounts.

In order to force a transition to LDAP/Active Directory only password authentication, the DCT administrator must either update the account records to remove the password, or disable local password authentication entirely.

6.5.3.2 Configure LDAP/Active Directory groups

In addition to being an authentication method, the LDAP/Active Directory integration can optionally also be used to retrieve additional attributes about the accounts authenticating: first name, last name, email address and group membership.

DCT only supports retrieving groups which are exposed as an attribute of the LDAP/Active Directory user record. DCT can not fetch groups membership from group records at the LDAP/Active Directory, and thus also does not support nested groups.

Group memberships are retrieved at authentication time, using the account credentials. DCT does not need credentials of an LDAP/Active Directory administrator, but will only be able to retrieve group memberships if LDAP/Active Directory users have the right to read the corresponding attribute.

This can be enabled by setting additional arguments to the domain API object.

search_base	<p>The Context name in which to search. Being specific enables faster LDAP search.</p> <p>To construct the search_base DN string according to your LDAP/Active Directory server, using an LDAP browser, navigate to a user, and then construct the search_base DN in reverse order from the User, up the folder hierarchy. For example:</p> <p>If a User DN is: CN=some-user-id,CN=Users,DC=mycompany,DC=co</p> <p>The corresponding search base might be: CN=Users,DC=mycompany,DC=co</p>
email_attr	<p>Name of the attribute in the LDAP/Active Directory server containing email addresses.</p> <p>Example: mail</p>
last_name_attr	<p>Name of the attribute in the LDAP/Active Directory server containing last names</p> <p>Example: sn</p>
first_name_attr	<p>Name of the attribute in the LDAP/Active Directory server containing first names</p> <p>Example: givenName</p>
group_attr	<p>Name of the attribute in the LDAP/Active Directory server containing group(s) membership. This can be a multi-valued attribute.</p> <p>Example: memberOf</p>
search_attr	<p>Name of the attribute in the LDAP/Active Directory server of which value corresponds to the username provided to the DCT login requests.</p> <p>For Active Directory, this is usually sAMAccountName.</p> <p>Example: If the search base is CN=Users,DC=mycompany,DC=co and the search_attr is principalName, DCT will search for a record with a principalName matching the username provided to the login request under the CN=Users,DC=mycompany,DC=co sub tree.</p>

object_class_attr	Restricts search to records with an objectClass matching this value. Example: person
-------------------	--

6.5.3.2.1 Active Directory example

The following requests enable LDAP authentication over SSL with an Active Directory server at address **activedirectory.company.co**, using the **us.company.co** domain, and configures optional attributes to retrieve first name, last name, email address, and group membership from the users sub-tree.

```
curl --location --request PUT 'https://<hostname>/v2/management/ldap-config' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "enabled": true,
  "auto_create_users": true,
  "hostname": "activedirectory.mycompany.co",
  "enable_ssl": true,
  "port": 636,
  "domains": [{
    "msad_domain_name": "mycompany.co",
    "search_base": "CN=Users,DC=mycompany,DC=co",
    "email_attr": "mail",
    "first_name_attr": "givenName",
    "last_name_attr": "sn",
    "group_attr": "memberOf",
    "object_class_attr": "person",
    "search_attr": "sAMAccountName"
  }]
}'
```

With the above config, when a user logs in with username John, DCT will:

1. Authenticate with the Active Directory server using the user principal name john@mycompany.co⁷⁴ and supplied password.
2. Search in the **CN=Users,DC=mycompany,DC=co** sub tree a record with **objectClass=person** and **sAMAccountName=john**.
3. Create or update a DCT Account record with the attributes extracted from the Active Directory server.
4. For each group membership found in the `memberOf` of the Active Directory server, an account tag is created with **key=login_groups** and value is the group name. These tags are protected (i.e cannot be modified within DCT) and can be securely used to control access groups membership.

⁷⁴ <mailto:john@mycompany.co>

As explained above, the **ldap-config/validate** API endpoint can be used to validate that each of the attributes corresponding to LDAP/Active Directory attributes.

6.5.4 SAML/SSO

The SAML 2.0 protocol allows DCT to delegate authentication to a SAML 2.0 compatible Identity Provider (Active directory federation services, Azure active directory, Ping federate, Okta, OneLogin, etc.). It only applies to web browser based interaction, and cannot be used for API access (scripting, integration).

Setting up SAML/SSO requires configuration changes both in the Identity Provider and DCT, so that trust can be established across both products.

When using SAML/SSO, DCT will uniquely identify accounts by email address, so make sure that records at the identity provider are configured with a unique email address.

DCT supports automatic account creation (or just in time account provisioning) when using SAML/SSO. When automatic account creation is enabled, accounts are created automatically when users login for the first time.

DCT allows group membership to be retrieved from the Identity Provider, which can be used to control access control authorization within DCT via DCT Access Groups. Using Identity Provider group membership allows DCT authorization to be managed per account group, and guarantees that authorizations in DCT reflect the organization structure which is expressed by group membership of the identity provider.

SAML/SSO is not mutually exclusive with other authentication methods, so enabling SAML/SSO is not disruptive (accounts configured with local password or LDAP/Active Directory authentication can still authenticate). In order to switch to SAML/SSO exclusively as authentication method for web browser interaction, perform the SAML/SSO configuration steps below and disable LDAP/Active Directory and Username/Password authentication. Note that API Key based authentication cannot be entirely disabled, but only administrators can create accounts with API keys.

6.5.4.1 Identity provider setup

Require that an administrator of the Identity provider used by your organization sets up a SAML 2.0 integration with DCT (an integration is sometimes called a Relying party trust, or an application).

The exact instructions are product specific, but the following input values must be provided:

Name	• Alternative name depending on product	Value

Single Sign-on URL	<ul style="list-style-type: none"> • SAML Assertion Consumer Service • ACS • Recipient URL • Destination URL • Relying party SAML 2.0 SSO • Service URL • Reply URL 	https://<dct-hostname>/v2/saml/SSO
Audience URI	<ul style="list-style-type: none"> • SP Entity ID • Relying Party trust identifier 	Any value can be selected, as long as the same value is set in the Identify Provider configuration and DCT configuration. We recommend: https://<dct-hostname>
Binding	<ul style="list-style-type: none"> • POST 	
Protocol		SAML 2.0 WebSSO protocol

The identity provider must be configured to include the email address as NameID attribute, and DCT will use the email attribute as a unique identifier for users when connecting via SAML/SSO.

6.5.4.2 DCT SAML/SSO setup

Once the configuration has been performed at the Identity provider, use the saml-config API endpoint to configure DCT accordingly. Copy the metadata from the Identity Provider using a web browser and provide it directly to DCT.

The Identity provider (IDP) metadata is a standardized XML document providing the SAML Service Provider (DCT) with the necessary information to verify the validity of incoming login requests and initiate a SAML/SSO login flow.


If `auto_create_users` is enabled, DCT will create accounts automatically when they login with SAML/SSO for the first time. If this is disabled, an administrator must create a DCT account with an email attribute matching the value from the SAML/SSO Identity provider before they can login. When `auto_create_users` is enabled, any user configured to authenticate via the Identity provider can authenticate to DCT, by default with an empty authorization set (i.e not being able to view any data or perform any action).

Without network access, provide the IDP metadata directly:

```
curl --location --request PUT 'https://<hostname>/v2/management/saml-config' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
```



```
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "enabled": true,
  "auto_create_users": true,
  "metadata": "<json-escaped-idp-metadata-xml-blob>",
}'
```

 The IDP metadata must be JSON escaped. On a terminal with `.jq`⁷⁵ installed, this can be achieved with the following command: `jq --slurp --raw-input <<< 'xml-metadata-here'`

6.5.4.3 Login

The SAML 2.0 protocol defines two login procedures: The Service Provider initiated flow starts by having users point their web browser to `https://<dct-hostname>/v2/saml/login` to login, while the Identity provider initiated flow starts at the Identity provider (details specific to Identity provider vendor). DCT supports both flows. The SAML/SSO authentication method is not intended for API interaction, and cannot be used with the Swagger UI.

After successful authentication, the web browser is redirected to the UI landing page and the the navigation bar can be used to go to the desired page. The session expires 24 hours after login.

6.5.4.4 Troubleshooting

There was an issue in SAML authentication: The assertion cannot be used before <timestamp>

The above error message, which is accompanied by `com.coveo.saml.SamlException`: The assertion cannot be used before <timestamp> error in the application logs, indicates that DCT was not able to validate the timestamp of the authentication provided by the Identity Provider. This is usually due to the system clock of the machine running DCT being incorrectly configured. Consider using NTP to maintain the machine's clock up to date.

There was an error fetching data

The above error message indicates that the current account does not have permission to view the data displayed on the page. Remember that, while DCT creates accounts automatically upon login when `auto_create_users` is enabled, by default accounts are created without any authorization and thus cannot see any data. Review the section below to see how SAML/SSO group membership can be assigned automatically at account creation.

⁷⁵ <https://stedolan.github.io/jq/>

6.5.4.5 Attributes mapping

As explained above, the only required attribute in the SAML Response (the message sent by the Identity Provider to DCT during login) is the **NameId** attribute which must be configured to a unique email address.

In addition to this, DCT allows for first name, last name, and group membership attributes to be included. The first and last names attributes will be stored as properties of the account object. For each group membership found in the SAML response attribute, an account tag is created with **key=login_groups** and value is the group name. These tags are protected (i.e cannot be modified within DCT) and can be securely used to control access groups membership.

In other to enable these optional attributes, update the Identity provider configuration to include them in the SAML response, and use the **saml-config** API endpoint to configure DCT with the name of the attributes configured in the Identity provider:

```
curl --location --request PUT 'https://<hostname>/v2/management/saml-config' \
--header 'Content-Type: application/json' \
--header 'Accept: application/json' \
--header 'Authorization: apk <your API key>' \
--data-raw '{
  "enabled": true,
  "auto_create_users": true,
  "metadata": "<json-escaped-idp-metadata-xml-blob>",
  "first_name_attr": "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/
givenname",
  "last_name_attr": "http://schemas.xmlsoap.org/ws/2005/05/identity/claims/
surname",
  "group_attr": "http://schemas.xmlsoap.org/claims/Group"
}'
```

With the above configuration, and a SAML Response as the following produced by the Identity Provider during login:

```

<?xml version="1.0" encoding="UTF-8"?>
<saml2:Assertion ID="id97923983167603821157180516" IssueInstant="2022-12-01T10:07:
12.856Z" Version="2.0"
  xmlns:saml2="urn:oasis:names:tc:SAML:2.0:assertion">
  <saml2:Issuer Format="urn:oasis:names:tc:SAML:2.0:nameid-format:entity">http://
/www.idp-demo.com/exk1fupjwz1YcMo290h8</saml2:Issuer>
  <saml2:Subject>
    <saml2:NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-
format:unspecified">john.doe@company.co</saml2:NameID>
    <saml2:SubjectConfirmation Method="urn:oasis:names:tc:SAML:2.0:cm:bearer">
      <saml2:SubjectConfirmationData NotOnOrAfter="2022-12-01T10:12:12.857Z"
Recipient="https://localhost/v2/saml/SSO"/>
    </saml2:SubjectConfirmation>
  </saml2:Subject>
  <saml2:Conditions NotBefore="2022-12-01T10:02:12.857Z" NotOnOrAfter="2022-12-0
1T10:12:12.857Z">
    <saml2:AudienceRestriction>
      <saml2:Audience>https://dct-demo.delphix.com</saml2:Audience>
    </saml2:AudienceRestriction>
  </saml2:Conditions>
  <saml2:AuthnStatement AuthnInstant="2022-12-01T10:05:07.916Z" SessionIndex="id
1669889232855.2084756273">
    <saml2:AuthnContext>
      <saml2:AuthnContextClassRef>urn:oasis:names:tc:SAML:2.0:ac:classes:Pas
swordProtectedTransport</saml2:AuthnContextClassRef>
    </saml2:AuthnContext>
  </saml2:AuthnStatement>
  <saml2:AttributeStatement>
    <saml2:Attribute Name="http://schemas.xmlsoap.org/ws/2005/05/identity/
claims/givenname" NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-
format:unspecified">
      <saml2:AttributeValue
        xmlns:xs="http://www.w3.org/2001/XMLSchema"
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs
:string">John
      </saml2:AttributeValue>
    </saml2:Attribute>
    <saml2:Attribute Name="http://schemas.xmlsoap.org/ws/2005/05/identity/
claims/surname" NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-
format:unspecified">
      <saml2:AttributeValue
        xmlns:xs="http://www.w3.org/2001/XMLSchema"
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs
:string">Doe
      </saml2:AttributeValue>
    </saml2:Attribute>
    <saml2:Attribute Name="http://schemas.xmlsoap.org/claims/Group"
NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified">
      <saml2:AttributeValue
        xmlns:xs="http://www.w3.org/2001/XMLSchema"

```

```

        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs
:string">Dev-Team
    </saml2:AttributeValue>
    <saml2:AttributeValue
        xmlns:xs="http://www.w3.org/2001/XMLSchema"
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" xsi:type="xs
:string">QA
    </saml2:AttributeValue>
</saml2:AttributeStatement>
</saml2:Assertion>

```

Would automatically create or update a DCT account with the following properties:

```

{
  "id": 94,
  "username": "john.doe@company.co",
  "firstName": "John",
  "lastName": "Doe",
  "email": "john.doe@company.co",
  "tags": [
    {
      "key": "login_groups",
      "value": "Dev-Team"
    },
    {
      "key": "login_groups",
      "value": "QA"
    }
  ]
}

```

6.6 Replace HTTPS certificate for DCT

By default, to enable HTTPS, DCT creates a unique self-signed certificate when starting for the first time. This certificate and private key are configured in the values.yaml file under:

```

proxy:
  crt:<certificate_value>
  key:<private_key_value>

```

To use your own certificates, these default values need to be replaced. They are Base64 encoded values of the certificate and key, respectively.

- To generate the Base64 encoded value of the certificate:

```
cat my_cert_file.pem | base64 -w 0
```

- To generate the Base64 encoded value of the key:

```
cat my_private_key.key | base64 -w 0
```

Generating a new TLS certificate and private key could require the assistance of your Security or IT departments. A new key pair (public and private key) will need to be created, in addition to a certificate signing request (CSR) for that key pair. Your IT department should be able to determine the correct certificate authority (CA) to sign the CSR and produce the new certificate. The common name of the certificate should match the fully qualified domain name (FQDN) of the host, as well as the FQDN as a Subject Alternative Name (SAN).

- After changing the **crt** and **key** values in values.yaml file. Run the HELM upgrade command:

```
helm upgrade -f values.yaml dct-services delphix-dct
```

- After upgrading the helm chart, restart the proxy pod to pick the new certificates:

```
kubectl rollout restart deployment proxy -n dct-services
```

6.7 External database support

6.7.1 Overview

DCT uses a PostgreSQL database to store all the persistent data powering the application (engines, VDBS, compliance jobs, accounts, permissions, etc.). By default, a PostgreSQL container image is packaged along with the DCT application and deployed along other pods onto the Kubernetes cluster, storing its persistent data into the `gwdatabase-data` persistent volume.

Alternatively, DCT can be configured to use an external PostgreSQL database instead, to which DCT connects over TCP and can then run anywhere (typically outside of the Kubernetes cluster).



External database support is only available for Kubernetes and OpenShift deployments. Docker Compose installations are not eligible.


6.7.2 Requirements


- **Database type:** PostgreSQL
- **Min supported version:** 12.16
- **Max supported version:** 14.10

6.7.3 Setup

6.7.3.1 PostgreSQL database setup

The following databases must be created prior to connecting DCT: **app**, **jobs**, **data-library**, **bookmarks** and **masking**.

 **masking** is newly added database for masking service and is applicable only for release 12.0.0 and above

 Quotes must be used to create or connect to the "**data-library**" databases, since its name contains a hyphen (-). Creating the database with a different name (such as **data_library**) will NOT work.

A PostgreSQL user must be created for DCT, with either the admin privilege or ALL privilege on the above databases. The following SQL script exemplifies the creation of required databases, granting privileges to a pre-existing `dct_user` (role) user.

```
CREATE DATABASE "bookmarks";
GRANT ALL PRIVILEGES ON DATABASE "bookmarks" TO dct_user;

CREATE DATABASE "data-library";
GRANT ALL PRIVILEGES ON DATABASE "data-library" TO dct_user;

CREATE DATABASE "jobs";
GRANT ALL PRIVILEGES ON DATABASE "jobs" TO dct_user;


CREATE DATABASE "app";
GRANT ALL PRIVILEGES ON DATABASE "app" TO dct_user;

CREATE DATABASE "masking";
GRANT ALL PRIVILEGES ON DATABASE "masking" TO dct_user;
```


6.7.3.2 DCT setup

Edit the `values.yaml` file to set `useExternalDatabase` to true and also provide the `dbHost`, `dbPort`, `dbUser`, and `dbPassword`, then run the following.

```
helm upgrade -f values.yaml dct-services delphix-dct
```

-  Previous versions of DCT required the aforementioned properties to be base64 encoded, but the values must be plain text values, as shown in the following excerpt:

```
useExternalDatabase: true
dbHost: "database-host.company.co"
dbPort: 5432
dbUser: "dct_user"
dbPassword: "dct_user_password"
```

-  If the password is stored as exemplified above in the values.yaml file, make sure to store this file in a secure location. Alternatively, set the value using the `--set` option flag in the HELM upgrade command, as shown in the following excerpt:

```
helm upgrade -f values.yaml --set dbPassword=dct_user_password dct-
services delphix-dct
```

- After upgrading the HELM chart, restart the pods to pick the changes:

```
kubectl rollout restart deployment data-library -n dct-services
kubectl rollout restart deployment data-bookmarks -n dct-services
kubectl rollout restart deployment jobs -n dct-services
kubectl rollout restart deployment gateway -n dct-services
kubectl rollout restart deployment masking -n dct-services
```

6.7.4 Backup and recovery

When using an external database, the `gwdatabase-data` persistent volume (created at default by DCT) to store database data is not used. Backing up and restoring the external database is not managed by DCT. Frequent or continuous backups are required, otherwise, DCT data will be lost.

The `gateway-data` persistent volume must still be backed up because it contains an encryption key, which is used to encrypt sensitive data at the application-level before being sent to the database. A backup of the external database cannot be restored successfully without a corresponding backup of `gateway-data`, as DCT would not be able to decrypt some of the data in the database. The encryption key in `gateway-data` does not change after having been initially created, so backups of it do not need to be scheduled at the same time as database backups.

6.7.5 External database migration or upgrade

The external database can be migrated to a different host, and/or upgraded to a different version at any time, as long the version requirements above are met. If the database is unavailable for a period of time, the DCT application will temporarily fail (internal server errors on all API calls), but will recover automatically without the need for a restart. However, in case of planned maintenance, upgrade, or migration, the following procedure should be followed:

1. Stop (shutdown) the DCT application.
2. Upgrade or migrate the database.
3. If necessary, set the updated database properties in values.yaml, and run `helm upgrade`.
4. Start the DCT application.

6.7.6 DCT upgrade

Before upgrading to a new DCT version, review the documentation to identify if the external database version is compatible. If the external database version is unknown, call the `metadata-database` API endpoint to get the information.

```
curl -k --location --request GET 'https://<dct-server>/v3/management/metadata-database' \
--header 'Accept: application/json' \
--header 'Authorization: apk <api-key>'

{
  "external": false,
  "version": "14.9",
  "database_product_name": "PostgreSQL",
  "major_version": 14,
  "minor_version": 9,
  "min_supported_major_version": 12,
  "min_supported_minor_version": 16,
  "max_supported_major_version": 14,
  "max_supported_minor_version": 9,
  "compatible": true
}
```

If the version of the external database is not compatible with the requirements of the DCT version being upgraded to, follow the instructions in the **External database migration or upgrade** section above before upgrading.



The selected PostgreSQL version to upgrade to must be compatible with both the currently running DCT version *and* the upgrade version.

After the upgrade, verify if the external database is compatible with DCT by inspecting the `compatible` property of the `metadata-database` API endpoint (as shown above).

6.8 DCT data backup, recovery, and migration

- i** This method is only applicable for **Kubernetes** and **OpenShift**.
- For Kubernetes, use the **kubectl** command prefix.
 - For OpenShift, use the **oc** command prefix.

The following directions will guide you through a Data Control Tower (DCT) back, recovery, and migration. They can be used for a “lift and shift” or a “blue-green” deployment. However, the backup and recovery steps are standardized and can be applied in one-off scenarios as well.

This page refers to the two DCT servers as **Initial** and **Destination**.

6.8.1 Example deployment scenarios

An example of a “lift and shift” deployment could be:

1. Backup the running *Initial* server pre-upgrade.
2. Upgrade the *Initial* server to the desired version and confirm functionality.
3. Backup the *Initial* server post-upgrade.
4. Install a new *Destination* server as the same *Initial* server’s post-upgrade version.
5. Restore the *Initial* post-upgrade backup to the *Destination* server.
6. Restart the *Destination* services and confirm functionality.
7. Shutdown the *Initial* server or maintain it for further testing.

An example of a “blue-green” deployment could be:

1. Backup the running *Initial* server.
2. *Install* a new *Destination* server with the same *Initial* server’s version.
3. Restore the *Initial* backup to the *Destination* server.
4. Restart the *Destination* services and confirm functionality.
5. Upgrade the *Destination* server to the newer version and confirm functionality.
6. Redirect traffic to the *Destination* server.
7. Shutdown the *Initial* server or maintain it for a future blue-green deployment.

6.8.2 Prerequisites

1. The *Initial* (source) DCT server is up and running.
 - a. Referred to as `svr_source` in any CLI commands.
2. The *Destination* (target) DCT server is installed in a separate Kubernetes cluster.
 - a. It is the same version as the *Initial* DCT server when the backup is taken.
 - b. Referred to as `svr_target` in any CLI commands.
3. We can share backup files from *Initial* to *Destination* environments.
4. Sufficient access to perform various `kubectl` commands on both the Initial and Destination clusters.

6.8.3 Directions

6.8.3.1 1. Backup

First, ensure the *Initial* DCT server is running, to take a successful backup. Then, run the following two commands:

```
kubectl cp <svr_source-gateway-pod>:/data data --namespace <svr_source-namespace>
```

```
kubectl cp <svr_source-masking-pod>:/data masking_data --namespace <svr_source-namespace>
```

```
kubectl exec -it <svr_source-database-pod> --namespace <svr_source-namespace> -- pg_dumpall -U postgres > postgres_db_all.sql
```

This will create three files: `data`, `masking_data` and `postgres_db_all.sql`:

- `data` is the Gateway pod's persistent volume containing various configuration information.
- `masking_data` is the masking pod's persistent volume containing various configuration information.
- `postgres_db_all.sql` is a complete database backup.

6.8.3.2 2. Restore

Ensure the *Destination* DCT server is running to restore successfully. In addition, make the `postgres_db_all.sql`, `data` and `masking_data` files available to the *Destination* cluster in the subsequent steps.

Then, run the following four commands:

```
kubectl cp data/encryption.key.sean <srv_target-namespace>/<srv_target-gateway-pod>:/
data --namespace dct-services
kubectl cp data/token_key_pair.sean <srv_target-namespace>/<srv_target-gateway-pod>:/
data --namespace dct-services
```

```
kubectl cp masking_data <srv_target-namespace>/<srv_target-masking-pod>:/data
```

```
kubectl cp postgres_db_all.sql <srv_target-namespace>/<srv_target-database-pod>:/tmp
```

```
for i in app bookmarks data-library jobs
{
  kubectl exec -it <srv_target-database-pod> --namespace <srv_target-namespace> --
psql -U postgres -c "drop database \"\$i\" with (FORCE)"
}
```

```
kubectl exec -it <srv_target-database-pod> --namespace <srv_target-namespace> -- psql
-U postgres -f /tmp/postgres_db_all.sql
```

Finally, delete and restart the DCT pods:

```
for i in `kubectl get pods --namespace <srv_target-namespace> | awk '{print $1}' |
grep -v jobs-cleaner | egrep "gateway|data-library|jobs|data-bookmarks|masking"`
{
  kubectl delete pod $i -n <srv_target-namespace>
}
```



After deleting the pods, Kubernetes will automatically recreate them and absorb the new database backup, and gateway volume data.

6.8.3.3 3. Additional environment configuration

The HELM chart's `values.yaml` contains information specific to your environment, such as certificates, hostname, or resource limits. You can update this information before or after the migration process. The standard installation and configuration process can be followed to update these values.

6.9 Exporting DCT logs to Splunk

6.9.1 Overview

This article provides some tips for configuring DCT (running on Kubernetes) to send logs to Splunk and extract useful information in Splunk.

6.9.2 Setting up a Splunk instance

Authenticate with Splunk via the web portal and install the third-party [Monitoring Kubernetes](#)⁷⁶ app directly via the Splunk UI, then enable HTTP Event Collector in Splunk and save the HTTP Event Collector token for future use.

6.9.3 Enable Splunk log forwarding

Once the Splunk instance is setup, follow the instructions to install Splunk logic in the Kubernetes cluster to forward logs to Splunk. This [blog post](#)⁷⁷ is a useful resources to understand the log collection and configuration options.

```
git clone https://github.com/splunk/splunk-connect-for-kubernetes.git
cd splunk-connect-for-kubernetes/helm-chart/splunk-connect-for-kubernetes
edit values.yaml
```

Edit `values.yaml`, at the minimum the host property (hostname of the Splunk collector) and token (of the HTTP Even Collector) must be set.

```
global:
  logLevel: info
  splunk:
    hec:
      # host is required and should be provided by user
      host: <insert-splunk-http-event-collector-hostname-here>
      # port to HEC, optional, default 8088
      port:
      # token is required and should be provided by user
```

⁷⁶ <https://splunkbase.splunk.com/app/3743>

⁷⁷ <https://faun.pub/logging-in-kubernetes-using-splunk-c2785948fdc0>

```
token: <insert-token-here>
```

Install the helm chart and after a few minutes DCT logs will be visible in Splunk.

```
helm install splunk-connect-for-kubernetes . -f values.yaml --set splunk-kubernetes-logging.fullnameOverride=splunk-logging
```

6.9.4 Search for events in Splunk

In the Splunk Cloud UI, via the “Monitoring Kubernetes” App, you can “search” for data sent by Kubernetes, as exemplified in the screenshot below. The bootstrap API key can be found as shown.

The screenshot shows the Splunk Cloud UI interface for the 'Monitoring Kubernetes' app. The search bar contains the text 'NEWLY' and the search results show a single event from 2/5/23 9:00:00.000 AM to 2/6/23 9:46:47.000 AM. The event details are as follows:

i	Time	Event
>	2/6/23 9:38:33.646 AM	NEWLY GENERATED API KEY: 1.r2hSuye1RLY86GF1q7nPH9ngQNo20rQLzVKzPZmsVTbdt44qpDz1c13N8XSW0f host = docker-desktop : source = /var/log/containers/gateway-77bdbc85bc-8htfh_apigw-services_gateway-640fb... : sourcetype = kube:container:gateway

The example screenshot below shows a search for “nginx”, with use of the “extract new fields” wizard on the bottom left, which has Splunk parse the Nginx access logs. A regexp is used to name some of the fields like ipaddress, endpoint, accountid, etc. The example runs a search to return API requests associated with accountid.

The screenshot shows the Splunk Cloud interface for Monitoring Kubernetes. The search bar contains the query: `index=*_ OR index=* sourcetype=kube:container:proxy AND account!="[:]" AND endpoint!="\POST /v2/management/accounts/search ?limit=50&sort=first_name HTTP/1.1\"200"`. The search results show 27 events from 2/5/23 10:00:00.000 AM to 2/6/23 10:09:16.000 AM. The events table is as follows:

Time	Host	Endpoint	Account
2/6/23 9:51:49.353 AM	docker-desktop	"POST /v2/management/accounts/5/tags HTTP/1.1"201	[4:admin-user-1]
2/6/23 9:51:42.976 AM	docker-desktop	"POST /v2/management/accounts HTTP/1.1"201	[4:admin-user-1]
2/6/23 9:51:35.547 AM	docker-desktop	"POST /v2/reporting/virtualization-storage-summary-report/search?limit=15&sort=-used_percentage HTTP/1.1"200	[4:admin-user-1]
2/6/23 9:51:35.547 AM	docker-desktop	"POST /v2/reporting/virtualization-storage-summary-report/search?limit=50&sort=-used_storage HTTP/1.1"200	[4:admin-user-1]
2/6/23 9:51:24.846 AM	docker-desktop	"POST /v3/access-groups/admin-user-1/policies HTTP/1.1"200	[1:-]

6.10 Generating a support bundle

6.10.1 Find the “collect_bundle.sh” script

- The support bundle tar file is available on the [downloads site](#)⁷⁸.
- Once the file is downloaded, untar the file to find the script.

```
dlpxuser@delphix:~/test$ tar -xzf dct-support-bundle-1.0.1.tar.gz
x ./
x ./collect_bundle.sh
x ./README
x ./VERSION
```

⁷⁸ <https://download.delphix.com/folder/1144/Delphix%20Product%20Releases/DCT>

6.10.2 Execute the “collect_bundle.sh” script when DCT is running in Kubernetes

- Transfer the "collect_bundle.sh" script to the machine where you have permissions to execute **kubectl** commands against the DCT pods.

 You must install the bash shell to generate a DCT support bundle, if it is not already.

- Execute the “collect_bundle.sh” script, which assumes a Kubernetes deployment by default. The script may need to run with "sudo", if root permissions are needed to run the **kubectl** commands.
- If the namespace is not the default "dct-services", use the "-n" flag and pass the correct namespace.

```
d1pxuser@delphix:~/test/tools/support-scripts/$ ./collect_bundle.sh -n
<custom_namespace>
....
DCT support bundle collection started at Thu Jun 22 12:35:05 EDT 2023
Collecting logs from all DCT containers...
....
```

6.10.3 Execute the “collect_bundle.sh” script when DCT is running in Docker-Compose

- Transfer the "collect_bundle.sh" script to the machine where you have permissions to execute **docker** commands against the DCT Docker-Compose application.

 You must install the bash shell to generate a DCT support bundle, if it is not already.

- Execute the “collect_bundle.sh” script with the "-d" parameter. The script may need to run with "sudo", if root permissions are needed to run the **docker** commands.

```
d1pxuser@delphix:~/test/tools/support-scripts/$ ./collect_bundle.sh -d
...
DCT support bundle collection started at Thu Jun 22 12:35:05 EDT 2023
Collecting logs from all DCT containers...
....
```

6.10.4 Find the generated support bundle tar file

The resulting support bundle will be located at `dct-support-****.tar.gz`, inside the current directory.

```
dlpxuser@delphix:~/test$ ls -ltr
total 316
-rw-r--r-- 1 65436 staff 104189 Feb 17 08:52 dct-support-<current_timestamp>.tar.gz
```

The support bundle tar file contains the following information:

- DCT logs for all of the containers.
- A java heap dump, `.hprof`, if one exists.
- A java thread dump and memory stats.
- The output of **docker stats**, if running in Docker-Compose.
- The output of **cpuinfo**, **meminfo**, and **mpstat** for each container, if running in Kubernetes.
- The output of **kubectl get pods -o json** for each container, if running in Kubernetes.



- The `collect_bundle.sh` generates a support bundle from a DCT engine running in Docker or Kubernetes.
- The resulting support bundle will be at `./dct-support-****.tar.gz` inside the container.
- The user must have privileges or permission to execute the **docker** or **kubectl**, commands in order to generate the support bundle.

7 Data governance

- [DCT administration](#) (see page 129)
- [Central governance workflows](#) (see page 162)
- [Insight reports](#) (see page 190)
- [Tabular customization support for DCT](#) (see page 192)

7.1 DCT administration

DCT delivers a management layer on top of all connected Delphix engines through surfacing object inventories, instrumenting all common Delphix operations, delivering a business metadata layer with tagging, and using those tags to drive attribute-based access control. This provides the ability for administrators to deliver a highly curated and secure Delphix experience for automation and end-users.

The image displays three panels illustrating DCT administration:

- Catalog:** Shows a list of databases and actions. The 'warehouse_Inventory' database is highlighted. Actions include Provision, Refresh, and Teardown. The 'SRE_Environment' environment is also highlighted. Below the list, it states: "Self-Service Access to Multicloud Data+Operations, Catalog, APIs, Integrations".
- Tagging:** Shows a list of data tags for 'Engineering Team Alpha'. Tags include Network (Non-Prod), Priority (High), Data Center (West), VDB Profile (Gold Copy), and Primary Owner (John Smith). Below the list, it states: "Tag Data for Control & Visibility".
- Attribute-based Access Control:** Shows a 'Central Permission' table for 'Engineering Team Alpha'. The table lists permissions: Create, Update, Delete, View, Refresh, and Rewind, each with a lock icon. Below the table, it states: "Protect Data with Global Attribute-Based User Access Management".


This section contains configurations handled under the **Admin page** in the DCT interface.

7.1.1 Operations in DCT


The Operations page is available under the **Admin** menu, showing a comprehensive list of all the DCT Jobs across Delphix infrastructures. It only shows operations that the user has access to.

The Operations Monitoring feature allows users, such as Database Administrators (DBA), Developers, CISOs, and Database Managers, to monitor system operations related to data virtualization and masking. Users can view all operations and their statuses (completed, in progress, or error) on the Operations page.

This page will display operations based on the user's **read permission** on the associated object. Click on a specific operation to view additional details, including the related object name, the operation type, the current status, a progress indicator (if in progress), and any additional metadata related to the operation.

 The Operations page is accessible from **Administration > Operations** or the **View All** button on the Operations Monitor Bar, which is described below.


Relevant details for the operation are shown, such as status with a graphical progress bar for running jobs, type of job, target id (refers to the id of the object on which the operation has been performed), start time, last updated time, etc. In versions 10.0.0 and above, a column shows the initiator of the operation.



Status	ID	Start Time	Last Updated Time	Target Name	Type	Initiator	
Started	9759fcaffe84e4cb3d...	Jan 18, 2024 4:32 PM ...	Jan 18, 2024 4:33 PM ...	as-win-src-si	Environment Refresh		Actions !
Completed	1660496c86f247cb1...	Jan 18, 2024 4:24 PM ...	Jan 18, 2024 4:29 PM ...	bwigt1	Environment Refresh		View >
Completed	214a7ab8f5d9420f99...	Jan 18, 2024 4:24 PM ...	Jan 18, 2024 4:25 PM ...	as-win-src-si	Environment Refresh		View >
Completed	e322dd537d1640958f...	Jan 18, 2024 4:20 PM ...	Jan 18, 2024 4:20 PM ...	DBOMSRB331B3	dSource Disable		View >
Completed	8c34a9878e3a41839...	Jan 18, 2024 4:13 PM ...	Jan 18, 2024 4:16 PM ...	VVSourceDB2019_UZ...	VDB Provision	Account-1	View >
Completed	79e32346f15b499ea1...	Jan 18, 2024 4:12 PM ...	Jan 18, 2024 4:12 PM ...	CDOMLOSr6004PDB...	VDB Delete	Account-1	View >
Failed	fa60dc504b7d4dc5b5...	Jan 18, 2024 4:11 PM ...	Jan 18, 2024 4:11 PM ...	CDOMLOSr6004PDB...	VDB Provision	Account-1	View >
Failed	c7c17c123cd449298...	Jan 18, 2024 4:07 PM ...	Jan 18, 2024 4:07 PM ...	DBOM_LUMW	VDB Provision	Account-1	View >
Pending	0a789ea40dec4f0e85...	Jan 18, 2024 3:37 PM ...		Book makr 2	Bookmark Create	Account-1	Actions !


The capture **above** represents an **admin user** view of the Operations page.

The capture **below** represents a **non-admin user** view of the same page.



Selecting a job will navigate to the Operations details page, which shows further details for that job, in addition to a progress bar. The details include target id, error, and warning logs. Continuous Data operations are the jobs initiated by the engine, as shown in the action sidebar.

The system will execute operations either manually by users or by automated tasks. If an operation encounters an error, the error details can be viewed on the Operation Details page. Notifications based on operation status changes will be covered in future brief and are not in the scope of this product brief.



7.1.1.1 Operations Monitor Bar

The Operations Monitor Bar appears at the bottom of pages to display user-initiated operations, with those completed being auto-dismissed. Operations with errors or in progress remain until the user dismisses them from the bar. In-progress operations display the operation type and object name. By default, the Operations Monitoring Bar toggle button will be available on all pages of the application.

When logging into the application, the Operations Monitoring Bar will be hidden by default. There will be two triggers for showing it:

- Click on the icon (Operations Monitor Bar toggle button) in the top right.
- Initiate an action like VDB refresh or compliance operation execution.

The bar has a **View All** button that navigates to the full Operations page and an actions menu, represented by a vertical "more" icon, that displays available actions for that operation. Visibility of the bar can be toggled on from the "memory" icon in the top right corner.

The Operations Monitoring Bar will show the following:

- Operations initiated by the user.
- Operations initiated during the current session.
 - Initially the Operations Monitoring Bar will always be empty upon login.
- Operations marked as **Monitor** by the user in the full Operations page.
- Operations currently running and failed.

7.1.1.1.1 Additional notes

- Close and reopen the Operations Monitoring Bar as needed by using the toggle button located in the top right corner (within the top banner).
- The successfully completed operations will auto dismiss. In other words, when any operation completes, the status changes from RUNNING to COMPLETED and it will be auto-dismissed/auto-removed from the Operations Monitor Bar. For any other status, such as CANCELED, ABANDONED, FAILED, SUSPENDED, TIMEOUT etc., auto dismiss will not occur.
- Click the three-dot menu icon on any operation in the Operations Monitor Bar to **dismiss** and remove any operation from the Operations Monitor Bar by selecting this action from the
- When logging out and logging in again, the Operations Monitoring Bar will revert to being empty. The application will not retain the state of the Operations Monitoring Bar between user sessions.
- The **View All** button will navigate users to the Operations page.
- The **View Details** button will navigate users to the Details page of the specific operation.

7.1.2 Tags

7.1.2.1 Tags management

DCT powers data governance with tags. These key-value pairs can be used to associate any business-level data with any Delphix object, to drive greater intelligence in automation, administrative workflows, data access, and reporting. Advanced search for tags is available.

Tags are individual attributes on every object exposed in DCT; from VDBs, to compliance jobs, and even users. There are no limits on tag count per object and character limits are set for flexibility to enable robust grouping.

7.1.2.2 Administrative tagging

Tags can be managed from the UI by selecting “View Tags” for a particular object on its global list page. The below example shows the tag configuration screen for a dSource “AGDatabaseSQL2016” and multiple tags have been added to characterize that particular object:

Add Tags for "AGDatabaseSQL2016".

App Team: Alpha ✕	🗑️ Remove
Application: Finance ✕	🗑️ Remove
Data Center: West Coast ✕	🗑️ Remove
Primary Owner: John Smith ✕	🗑️ Remove
Secondary Owner: Jane Doe ✕	🗑️ Remove

<input type="text" value="Enter Tag Key"/>	<input type="text" value="Enter Tag Value"/>	+ Tag
--	--	--------------------

Cancel Add Tags

DCT tags enable complex searching to enable intelligent reports. A demonstration using the above example dSource and using expression-based search to filter dSources with the {App Team: Alpha} tag.

dSources
An overview of data sources across your Delphix infrastructure.

Search: tags CONTAINS {key EQ 'App Team' AND value EQ 'Alpha'}

Expression Search
Delphix supports expression-based searching. See below for examples of the search expressions.

Status	Name ↑	Type	Engine	Tags	
UNKNOWN	AGDatabaseSQL2016	MSSql	amit-engine	View tags (5) >	View >
UNKNOWN	AppFS_cust_master	Unstructured Files	Test Engine	View tags (1) >	View >
UNKNOWN	AppFS_master	Unstructured Files	Test Engine	View tags (1) >	View >

View Tags for "AGDatabaseSQL2016".

- App Team: Alpha
- Application: Finance
- Data Center: West Coast
- Primary Owner: John Smith
- Secondary Owner: Jane Doe

© 2022 Delphix. All Rights Reserved. Private and Confidential

7.1.2.3 Tags powering attribute-based Access Control

Tags also power the DCT permissions system for both Accounts (users) and Role Scopes (object entitlements). The below example shows an Access Group (Alpha Team) with the Accounts tab on display. Notice that the accounts tab has {App Team: Alpha} under “tag mapping”, which automatically attributes any users with the {App Team: Alpha} tag.

DATA CONTROL TOWER | Data | Compliance | Insights | Admin

Access Group: **Alpha Team**
e35e45ad-4c19-41fb-a8c0-b8a0df9e68e9

Accounts | Roles | Access

Tag Mapping [Edit](#)

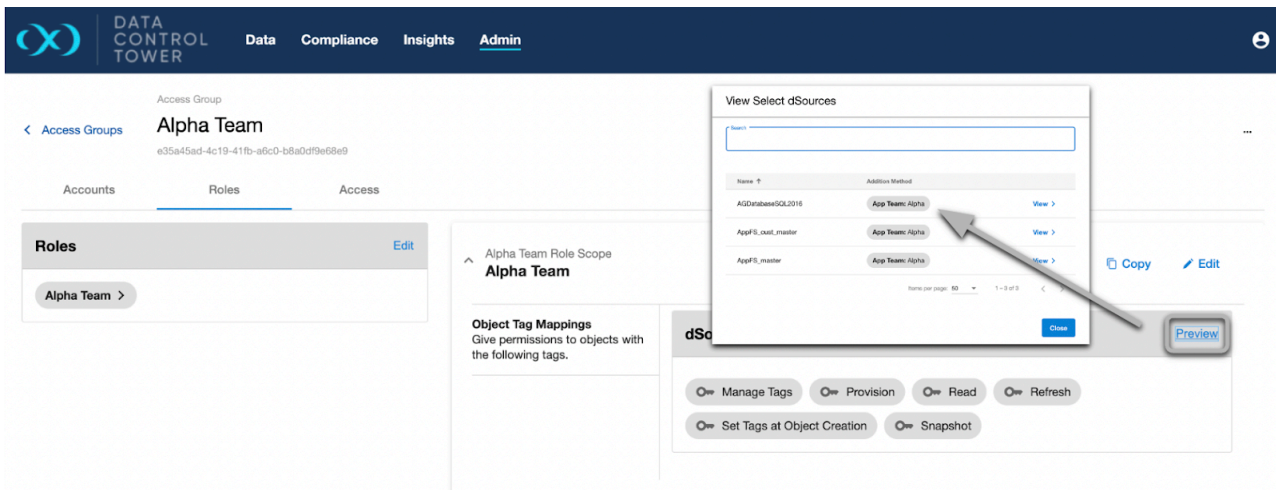
- App Team: Alpha

Associated Accounts [+ Manually Add Accounts](#)

Search: Search (All) ▾

ID	Username ↑	Name	Email	Addition
4	John Smith	John Smith	johnsmith@company.com	App Team: Alpha View >

The same goes for Scoped Roles under the “Roles” tab. The Alpha Team role has been mapped to the {App Team: Alpha} tag and all dSources with that same tag are automatically attributed.




7.1.3 Access Control

Data Control Tower (DCT) fundamentally changes how application teams are governed across the Delphix Platform to ease expansion and management burden. Previously, Delphix administrators were focused on managing individual user-level access on each engine. This made it difficult as teams increased their data set requirements. This inevitably led to more time managing engine access and not rolling out test data management (TDM) practices. Now with DCT, all users are managed and access their data sets through a centralized server. This makes it easier for administrators to manage the Delphix Platform and application teams to utilize the self-service capabilities.

To take advantage of DCT’s new capabilities, Delphix administrators will implement a centralized Attribute Based Access Control (ABAC) model. This is performed by consolidating permission management from the engines to DCT, implementing Access Group policies, and assigning Object tags. The flexibility of this approach ensures your company’s required security model can be maintained or even further refined.

The below picture attempts to show the shift in access models. In the original Engine Model, the engines were isolated from one another. No access control mechanisms were shared between Engines. In the DCT Model, Delphix administrators will manage applications teams directly through DCT. Those application teams will log directly into DCT. Only administrators will log into the Engines for advanced usage.




7.1.3.1 Access model overview

Data Control Tower implements a model that you might find in other types of software called Attribute Based Access Control (ABAC). This model is incredibly flexible but requires detailed configuration to perfect your use cases. In our model, there are four entity types which are defined below. Understand each entity as they are the foundational blocks of DCT's ABAC model.

Entity	Description	Managed By
Accounts (aka Users)	A single or shared user who can authenticate with DCT (UI or API).	Create manually or via Identity Provider (IdP), such as SSO or LDAP. Accounts are independent of Delphix Engines.
Access Groups	A collection of accounts that share one or more characteristics, such as a Team or Permission set. Equivalent to an Active Directory group.	Manually created. Populated manually or via the 'login_groups' tag.
Roles and Permissions	The collection of read, write, and delete permissions forms a reusable, named role.	Some roles are provided out of the box, but Admins can build their own from the available permissions. Individual permissions are immutable.
Objects	Units, such as VDBs, Bookmarks, and Environments, that are managed across the Delphix Platform.	Automatically identified by DCT from the connected engines. Assigned to Roles via various models. The CD and CC Engines supply these objects.

Each entity is linked to another through manual or automated assignment. A manual (or direct) assignment is a good approach for early implementations. However, it can be challenging to maintain as teams grow. As an alternative, Tagging is suggested as it performs automatic assignments based on your custom configuration. The below diagram shows how each entity is linked together. The directions below start with Accounts creation to Access Groups with Role assignments and finish with Object mappings.

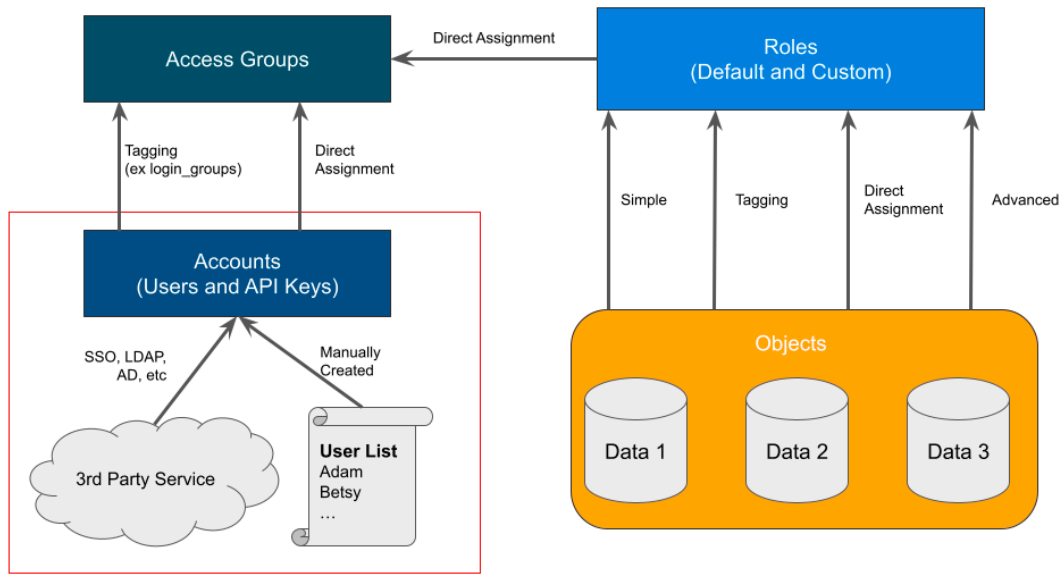


Understanding your team structure is imperative to identify the best access model. Usually, organizations have existing groupings defined in their Identify Provider (IdP). These groups are typically organized in one of two ways (a) a team dedicated towards a central goal (such as a product development team) or (b) a group of individuals with similar permissions (such as Security Administrators). Understanding the purpose of each group should be a guide in how the Roles and Permissions are designed. For example, the Alpha product development team might have full permission to manage existing VDBs and create new bookmarks for their team's "Alpha" objects. On the other hand, Security Admins might have sweeping read and disable access across the entire platform to ensure compliancy. Iterating through each Access Group and designing custom, but re-useable roles, based on the [Principle of Least Privilege](https://en.wikipedia.org/wiki/Principle_of_least_privilege)⁷⁹, will produce a streamlined rollout.

7.1.3.2 Accounts: Manual, LDAP/AD, or SSO/SAML

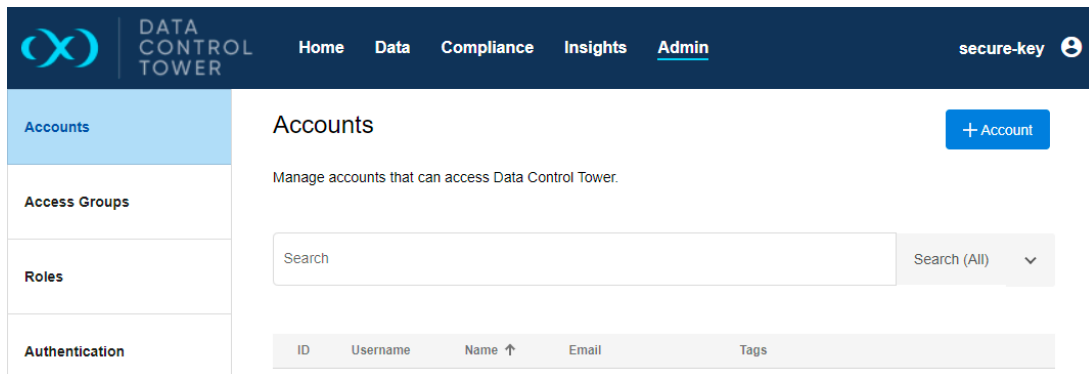
Goal: Import or create user accounts. Complete either the Manual or LDAP/SOO configuration.

⁷⁹ https://en.wikipedia.org/wiki/Principle_of_least_privilege



Manual (User List)

Navigate to **Admin > Accounts**, click the “+ Account” button, and complete the form.



Manual accounts are great for testing user access or providing a service account. Take note of the checkboxes by which you want this user to access DCT.

Create Account

Create an account and specify the authentication method.

Username

First Name Last Name

Email

Add "admin" role to account access group

Enable Authentication with username and password

Generate API Key

Add Account Tags
Account tags help you group and organize DCT accounts

Enter Tag Key Enter Tag Value


When you have specified all required values, select the “Create Account” button. By default, this user will have no permissions.

LDAP/Active Directory or SSO/SAML (3rd Party Service)

Navigate to **Admin > Authentication**, click “Edit” for either LDAP/AD or SSO/SAML, and complete the form. Ensure “Auto-create Users` is enabled. It can be disabled at any time.

If you need guidance on how to configure, follow the directions here:

- [LDAP/AD Directions \(see page 103\)](#)
- [SSO/SAML Directions \(see page 111\)](#)



Accounts

Access Groups

Roles

Authentication

Authentication Settings

LDAP/AD Edit

Status
Disabled

LDAP Server
None

Port
None

Auto-create Users
Enabled

Enable SSL
Enabled

Insecure SSL
Disabled

Unsafe SSL Hostname Check
Disabled

SSO/SAML Edit

Status
Disabled

Entity ID
None

First Name Attribute
None

Last Name Attribute
None

Group Attribute
None

Response Skew
120

Auto-create Users
false

Once configured, Accounts will be automatically created when a user successfully logs in.




This is functionally different from the old Engine model. Previously, the Account was created manually before they could log in.

Recommended: LDAP/Active Directory Domains

It is highly recommended that we also configure group membership during this stage. By defining the metadata attributes in the option Domain fields, DCT can automatically assign users to Access Groups. If configured correctly, you will see an automatically generated `login_groups` tag on recently logged-in accounts. If an Account does not have the tag, then (a) the Domain configuration is invalid, or (b) they should re-login. The `login_groups` tag is the only tag that cannot be specified on an Account manually.

[LDAP/Active Directory Domain Groups Directions \(see page 108\)](#)




Example Account with the login_groups tag.

Please test these new accounts out by logging in on another browser. By default, these accounts will not have any permissions and not see anything. In the following steps, we will give them access. In addition, if configured, we'll take advantage of login_groups or a custom tag.

7.1.3.3 Access Groups: Creation and account assignment

7.1.3.3.1 Goal: Create an Access Group and assign Accounts directly or through Tags.



Access Group Creation

Next, navigate to the **Admin > Access** Groups tab, select the “+ Access Group” button, and complete the presented form. As described previously, these groups are based on existing teams or users with similar access. If you successfully configured the Active Directory’s Domain Groups, you can specify the `login_groups` tag and value here. Or specify a custom Tag, such as “Team: Alpha”.

You can also select Roles if you already know which should be applied. Otherwise, ignore it.

Submit the form once you are happy with your new group.



Unlike an Account, you can specify the `login_groups` tag on an Access Group shown in the picture above.

On completion, you will be presented with a page similar to the one below. Unfortunately, it’s empty. Let’s add some associated Accounts now.


Manual (Direct) Assignment

Select the “+ Manually Add Accounts” button, select the desired Account, and then “Add Account”. Immediately, you’ll see it presented in the Associated Accounts list.

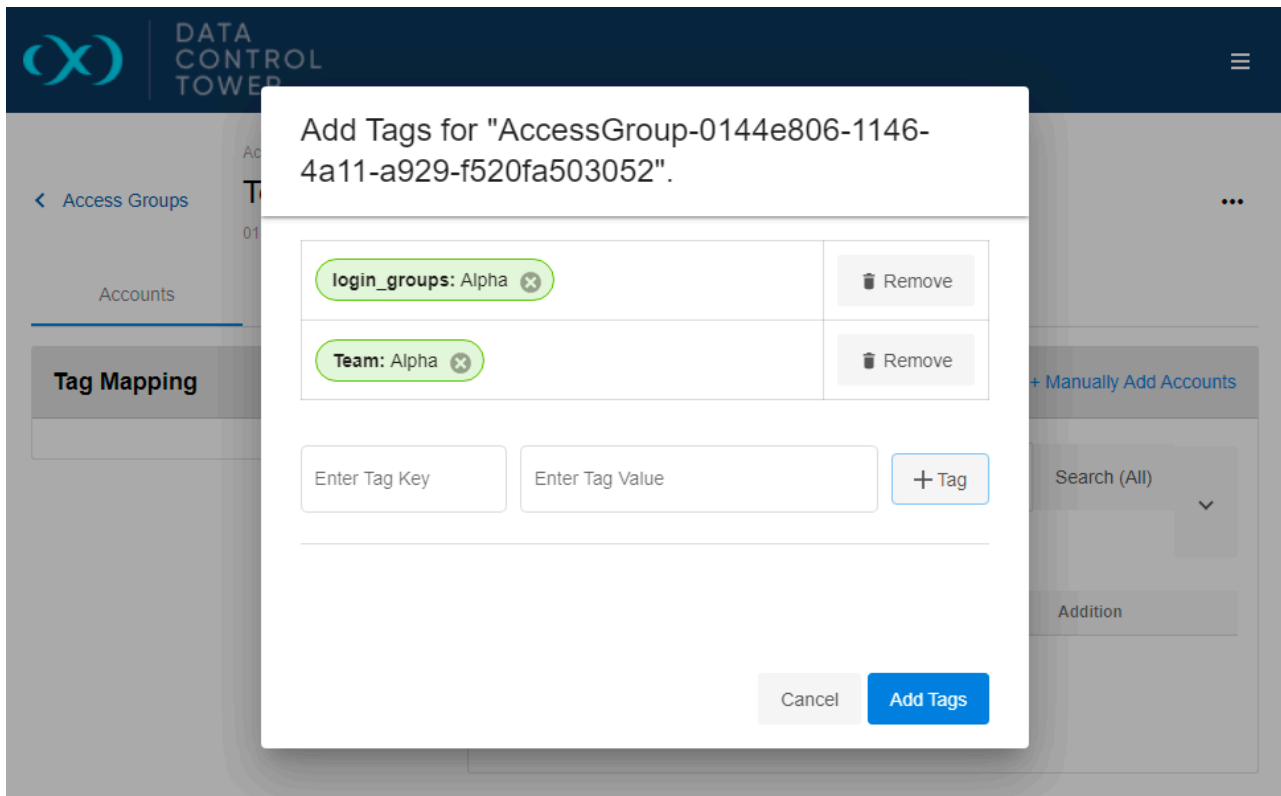
This is a good solution for quick management. However, it can be cumbersome as usage grows. Therefore, we recommend tags!

Tag Assignment

First, navigate to the **Admin > Accounts** tab and select an existing Account. (Feel free to create another one!) Once selected, add a custom Tag such as "Team: Alpha". If one already exists on the Account, such as "login_groups", remember it.



Next, navigate back to the Access Group, select the Tag Mapping's "Edit" button, and specify that same Key: Value pair. It might look similar to the below picture.



In this example, the “Team: Alpha” and “login_groups: Alpha” were added through the Access Group’s Tag Mapping widget. If configured successfully, your Access Group might look similar to the below picture. If you remove the Access Group or Account’s tag, you will see Account automatically removed from this listing.

i The “login_groups” tag functions identically to a custom tag within the Access Group. Again, the only difference is that it’s automatically assigned to the Account.

Access Group

< Access Groups **Team Alpha** ...

0144e806-1146-4a11-a929-f520fa503052

Accounts Roles Access

Tag Mapping [Edit](#)

login_groups: Alpha

Team: Alpha


Associated Accounts [+ Manually Add Accounts](#)

Search Search (All) ▾

ID	Username	↑ Name	Email	Addition
2	secure-key			Team: Alpha View >
39	test	Test User	user@test.com	login_groupsAlpha View >


This section taught us how to organize Accounts into different groups. This allows us to keep permission sets separated. Feel free to experiment with new Access Groups, Tags, and Accounts. If you still need additional pointers, review our [Access Groups Documentation](#) (see page 141).

7.1.3.4 Roles: Creation and assignment




Role Investigation and Creation

Navigate to the **Admin > Roles** tab. Here we see a list of DCT’s default Roles. Each role has its selection of Permissions, such as Read VDB, Delete Bookmarks, Modify dSources, etc. Select “View” on the “devops” role to see its permissions.




On the left-hand side, you can see a description, the Access Groups it’s currently a part of, and any assigned Tags. On the right-hand side, is the complete list of permissions. For example, you can see here that the “devops” role has “Manage Tags” and “Read” permissions on the CDBs objects. These various permissions make up the Role’s identity.

 DCT's default roles are immutable.

Role Creation

Now we understand what it's composed of, let's create one. Navigate back to the Admin > Roles tab and select the "+ Role" button. Give the Role a custom name, sample description, and add all the permissions you want. In my simple example, I gave it the "VDBs > Read, Refresh, and Manage Tags" permissions. If you need to grant permission for the entire category, select the header checkbox, such as "Access Groups" or "Bookmarks". If you only want a portion of that Object group, then click the little arrow icon to open up the complete set of options and select the targeted permissions.



The screenshot shows the 'Create Role' dialog box in the Data Control Tower interface. The dialog has a title bar 'Create Role' and a close button in the top right. It contains the following fields and options:


- Name:** A text input field containing 'Limited Monitoring'.
- Description:** A text area containing 'Simple Blog Post Role'.
- Permissions:** A list of checkboxes under the 'VDBs' category:
 - Access Groups
 - Accounts
 - VDBs
 - Read
 - Provision
 - Stop
 - Start
 - Disable
 - Create
 - Refresh
 - Manage Tags

At the bottom of the dialog are 'Cancel' and 'Create' buttons.

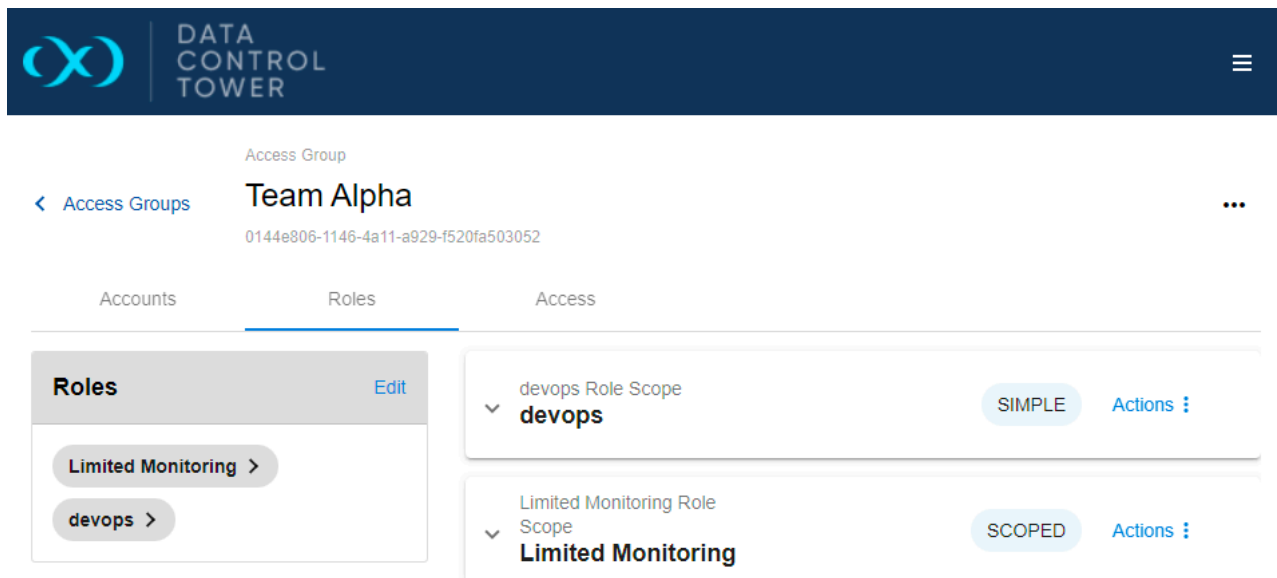
Once happy with your selection, click “Create”. You can modify your Permissions further on the presented page.

Role Assignment

Roles, by themselves, provide no access. You must first assign them to an Access Group and a set of Objects before their permissions are applied to an Account. Let’s do the first part now. Navigate back to the **Admin > Access** Groups tab and “View” your previously created Access Group. Select the “Roles” subtab and then “Edit” within the Roles widget.




Now you can assign default Roles, such as “devops”, and your newly created Role. You might recall that Role assignment was also possible during Access Group creation. On Save, your Access Group might look like the following.



Immediately on assignment, all users within the Access Group will now have the permissions assigned to them through these roles. (Since you are currently an Admin user, you must log in as your test Account user.) However, you might notice that this user has full access to every object on DCT. The following section will define the Role scoping modes and refine the Account Object access.

7.1.3.5 Objects: Refine permission to targeted objects



Every Access Group's Role has its own set of Objects to which the permissions are applied. In the previous section, we defined the permissions, and now we select the Objects. Objects are assigned in three different modes. They are listed below with their method of application:

1. Simple - All Objects within DCT.
2. Scoped
 - a. Tags - Objects with matching Tags.
 - b. Direct - Objects manually assigned.
3. Advanced Scoped - Objects are assigned directly on the permission action (such as Read Bookmark, Edit Bookmark, or Delete Bookmark) using Tags or Direct Assignment.

Edit devops Role Scope



- Scope Mode
- Add Tag Mappings
- Add Objects

Scope Mode

Select the mode that will determine what objects are affected by the role permissions.

Scope Name
devops

- Simple
The permissions apply to every object applicable to the permissions found in the Role.
- Scoped
The permissions is set to objects that either match the mapped tags or have been manually added.
- Advanced Scoped
The permission scope is set for each permission and gives the maximum level of granularity.

Cancel Back **Next** Submit

We will work through the first two in this post, Simple and Scoped. Advanced is easier to comprehend afterward and a solid self-lead challenge. Before diving into this section, we recommend that your DCT server has a handful of objects, such as Bookmarks or VDBs.

Simple

If you have been following the post steadily, you should have two Roles assigned to your Access Group. In my example I have “devops” and “Limited Monitoring”. Both are given the “Simple” mode by default. We can see the breadth at which this Role governs by selecting anywhere on its row and then the “Preview” button on the right-hand side.

View Select Bookmarks


Name ↑	Addition Method
Bookmark Test 123	All Bookmarks
CRIT-BUG-2	All Bookmarks
CRIT-BUG-4	All Bookmarks
CRIT-BUG-5	All Bookmarks

In this example, we select the “Preview” list for Bookmarks. It displays every Bookmark this role has access to and the method to which they are applied. If we wanted to validate, we could log in as a user on this Access Group and verify the permissions are applied. However, this is an easier way for Administrators to confirm without switching logins. Because this is a “Simple” scope, every object is available, so this view is not particularly intriguing. In the next part, we’ll refine our Role.

i If you do not see any objects listed in the “Preview” widget, that object might not be available to DCT. This could be because (a) engines are not connected, (b) the DCT-only object (such as Bookmarks or VDB Groups) is not created, or (c) permissions are being enforced correctly.

Scoped - Direct

Let’s change the mode to “Scoped” and target a subset of VDBs. On the Access Group > Roles tab, select the **Action > Edit** button of your chosen Role.



A new wizard will appear with the Simple, Scoped, and Advanced Scoped options. Change the Role’s mode from “Simple” to “Scoped”. Skip the “Add Tag Mappings” for now and select “Next” to move to “Add Objects”. You will be presented with a long list of the objects available to DCT. This is where you can manually assign specific DCT objects.

Edit devops Role Scope ✕

Scope Mode

Add Tag Mappings

Add Objects

Add Objects

Manually select objects to include in addition to the objects being added dynamically from tag mappings.

Bookmarks

CDBs

Database Templates

Dataset Groups

dSources

Engines

Environments

Sources

VCDBs

VDBs

No Manually Added VDBs

Manually add VDBs

Name
<input checked="" type="checkbox"/> Mysql_MASK
<input checked="" type="checkbox"/> MySQL_DEV
<input checked="" type="checkbox"/> Mysql_QA
<input type="checkbox"/> Postgres_MASK
<input type="checkbox"/> Postgres_DEV
<input type="checkbox"/> Postgres_QA
<input type="checkbox"/> AppFS_MASK
<input type="checkbox"/> AppFS_DEV
<input type="checkbox"/> AppFS_QA

Cancel
Back
Next
Submit

Scroll down the Object type list and select VDBs. Next, choose the “Manually add VDBs” radio button and, on the right-hand list, select a couple of VDBs. Feel free to add other available Objects too. When happy with your selection, press the “Submit” button. This set of actions should change the Role’s “Simple” mode to “Scoped” mode. Let’s verify by, again, opening the Role’s row, scrolling to your chosen Object Type, and selecting the VDBs’ “Preview” button.

The screenshot shows the 'View Select VDBs' modal window in the Data Control Tower interface. The modal contains a search bar at the top, followed by a table with the following data:

Name ↑	Addition Method	
MySQL_DEV	Manual	View >
Mysql_MASK	Manual	View >
Mysql_QA	Manual	View >

Below the table, there are pagination controls: 'Items per page: 50' and '1 - 3 of 3'. A 'Close' button is located at the bottom right of the modal. In the background, a 'Preview' button is highlighted with a red box.

In my example, the same three VDBs I selected during permission configuration are shown here. If you want to verify manually, log in as your other test user and confirm.

i Any other Roles or Access Groups assigned to this user might affect its visibility. So if you do this test, ensure it's not accidentally pulling in another permission set.

Scoped - Tags

Direct Assignment is a solid strategy for early onboarding and one-off requirements. However, as we expand our consumption of Delphix, I suggest leveraging the Tagging mechanism to assign permissions quickly. Similar to the Account & Access Group's "login_groups" tag, we can assign tags to Objects and Roles to immediately grant or restrict access. This is the recommended approach for a robust production implementation.

Before jumping back into a Role, navigate to the top-level Data >VDBs tab. (If you don't have any VDBs, select another tab with available objects.) Here identify a test object and select the "Add Tags" button.

The screenshot shows the Data Control Tower interface with a modal dialog titled "Add Tags for 'AppFS_cust_QA'". The dialog contains a list of existing tags, including "Team: Alpha" with a remove button. Below this is a form to add a new tag, with "Environment" entered in the "Enter Tag Key" field and "Dev" in the "Enter Tag Value" field. A red box highlights the "+ Tag" button. At the bottom of the dialog are "Cancel" and "Save" buttons. In the background, a table lists objects with "Add Tags" buttons, one of which is also highlighted with a red box. A red arrow points from the "+ Tag" button in the dialog to the "Add Tags" button in the table.

In this form, we assign a simple Key-Value pair. This pair helps govern access and maintain the organization of the Delphix Platform. I've selected the "Team: Alpha" and "Environment-Dev" pairs in my example. Repeat the process for a couple of other objects using similar or different Key-Value pairs. As I explained earlier in this post, we can define and create an organizational structure in many ways. If you prefer other pairings, please experiment, such as with Geography, Age, DB Type, or Importance.

Next, let's take advantage of the created tags in the Access Model. Navigate back to your test Access Group, select the Roles tab, and edit the Role we modified previously. Because "Scoped" is already chosen, press the "Next" button, but this time stop on the "Add Tag Mappings" view. Similar to your Object's tag assignment, specify one or two of the same Key-Value pairs here.

Edit devops Role Scope



- Scope Mode
- **Add Tag Mappings**
- Add Objects

Add Tag Mappings

Select tags that will dynamically add objects that are assigned the mapped tag. In Advanced Mode, you can select tags for each permission included in the role.

Team: Alpha ✕	Remove
----------------------	--------

Enter Tag Key	Enter Tag Value	+ Tag
---------------	-----------------	-------

Cancel Back **Next** Submit

In my example, this process will assign the Objects with the chosen "Team: Alpha" tag to this "devops" Role. Thus, granting the set of permissions defined by "devops". Finally, we can verify again by completing the form and previewing the Role.

The screenshot shows the 'View Select VDBs' modal in the Data Control Tower Admin interface. The modal contains a search bar and a table of VDBs. The table has two columns: 'Name' and 'Addition Method'. The rows are as follows:

Name ↑	Addition Method
AppFS_cust_DEV	Team: Alpha
AppFS_cust_MASK	Team: Alpha
AppFS_cust_QA	Team: Alpha
MySQL_DEV	Manual
Mysql_MASK	Manual
Mysql_QA	Manual
OracleDEV_CHT	Team: Alpha

At the bottom of the table, there is a pagination control showing 'Items per page: 50' and '1 – 9 of 9'. A 'Close' button is located at the bottom right of the modal.

In my example, we can see a mix of objects assigned to this role through Tags and Manual (direct) assignments.

At this point, challenge yourself by adding and removing tags to different Roles and Objects to understand the flexibility of the ABAC model. [If you need a deeper dive into Tags, read our documentation here.](#) (see page 132)

7.1.4 VDB templates



For additional detail on VDB templates, visit the “Configuration Settings for Oracle VDBs” article in the Continuous Data Engine documentation.

DCT has implemented a global VDB template system to centrally manage and apply VDB templates for any and all VDB provisioning workloads. This feature works as an extension of the local VDB template system on Continuous Data Engines as a means of enforcing VDB configuration standards and policies uniformly.

DCT Admins have the choice of either importing pre-existing VDB templates from a local engine or creating net-new templates from within DCT.

7.1.4.1 Creating templates

Users can create Database Templates directly via DCT, which can then be used on VDBs across their engines. The DCT API interface for creating templates is equivalent to that of on-engines, requiring a name and sourceType, and optionally taking in a description and the list of config parameters. Here's a sample CURL command:

```
curl -X 'POST' \
  '<https://<APPLIANCE_ADDRESS>/v2/database-templates' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "name": "vdb-config-template-1",
    "source_type": "OracleVirtualSource"
    "parameters": {"config1": "value1", "config2": "value2"}'}
```

This will result in a new DCT **DatabaseTemplate** object, which can then be viewed using the **List/Get/Search** APIs.

7.1.4.2 Importing templates

Unlike many other Delphix objects, DCT is not automatically pulling in all the Database Templates from registered engines and creating DCT objects out of them. It is often the case that users have already made arrangements and have copies of their templates across their engines. DCT does not blindly import the templates to avoid generating duplicates, leading users to consolidating and clean up. Instead, DCT provides an import API that can be used to selectively choose which engines they wish to import their templates from, along with an API to undo imports. The import workflow has a couple of things to be aware of:

- The user cannot be selective of which individual templates to import from an engine. The import API will pull ALL templates from that engine.
- Import is allowed only one time per Engine. After an initial import, subsequent imports will be blocked, and it is assumed that a user will use the DCT APIs to create more templates.
- In the event that an import was done on accident or no longer desired, the undo import API can be called to delete all the imported templates from the selected engine. This will result in the removal of all DCT Database Templates that were created as a result of the import.
- If an imported template is later used on a VDB running on a different engine than where it was originally imported from, then the undo import flow is also prohibited, as DCT can no longer safely delete a template that is in use elsewhere.

Import templates from the engine:


```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v2/database-templates/import' \
  -H 'accept: */*' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "engine_id": "3"
  }'
```

Undo the imported templates from engine:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v2/database-templates/undo-import' \
  -H 'accept: */*' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "engine_id": "3"
  }'
```

7.1.4.3 Using templates

DCT Database Templates can be used by specifying the `template_id` property at VDB provisioning time, or by updating the `template_id` on an existing VDB. In either case, DCT will deploy the template to the respective engine and bind the template with the VDB. When a DCT Database Template currently in use is updated or deleted, those changes are propagated to the respective VDBs and engines.

-  If a VDB has the same parameter called out in both VDB template and individual setting, the value specified in the template will take precedence. The individual parameter value will only be used if the VDB template is removed.

Updating a VDB to use `template_id`:

```
curl -X 'PATCH' \
  'https://<APPLIANCE_ADDRESS>/v2/vdbs/1-ORACLE_DB_CONTAINER-1' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "template_id": "319db966-961c-4977-a444-14d337aa3276"
  }'
```



Updates to a VDB template will propagate to all associated VDBs.

7.1.5 API metering

7.1.5.1 API metering instructions

DCT employs a per API consumption model, which requires API metering and periodic reporting to Delphix Customer Success. To support reporting of API consumption, DCT offers an API consumption reporting endpoint called, “api-usage-report”. This report will provide a list of all unique API endpoints and how often they were used over the specified time period sorted by API and method.

Required inputs

- File type: CSV or JSON (CSV file types are compatible with most spreadsheet-style software like Excel or Google Sheets)
- Start/end date: The default start date is “when DCT was installed” and the default end date is the “time when the report was generated”.

7.1.5.1.1 Example cURL call

```
curl --location --request GET 'https://[Inser_DCT_Server]/v2/reporting/api-usage-report/?end_date=2022-06-14T09:00-04:00&start_date=2022-06-01T00:00Z' \
--header 'Content-Type: application/json' \
--header 'Accept: text/csv' \
--header 'Authorization: apk 1.xxxxxxxx'
```

Example output

```
api_endpoint,api_method,api_count
"/v2/management/api-clients",GET,2
/v2/management/engines,GET,1
"/v2/management/engines/search",POST,1
"/v2/reporting/api-usage-report",GET,2
```

7.1.6 Client telemetry

DCT provides complete flexibility to clients on how to attribute their API calls. DCT captures the value provided in an optional HTTP header (`X-Dct-Client-Name`) and standard, mandatory HTTP header (`User-Agent`) for the purpose of attributing an API call. These values are stored as `client_name` and

`user_agent` in the backend, and can be queried in the report. Below are some examples of how this can be used.

Example one

Clients can view the report grouped on the basis of client name and API classification by providing a `group_by` query parameter.

```
curl --location 'https://[Inser_DCT_Server]/v3/reporting/api-usage-report?
group_by=client_name%2Ckind' \
--header 'Authorization: <api_key>' \
--header 'Accept: text/csv'

api_endpoint,api_method,api_count,kind,client_name,user_agent,dct_version
,,2,management,client-1,,
,,5,management,client-2,,
,,20,management,client-3,,
```

The `group_by` parameter supports any combination of properties from `api_endpoint`, `api_method`, `kind`, `client_name`, `user_agent`, and `dct_version`. All properties function as their name describes, where `kind` corresponds to API classification.

Example two

Clients filter the records for a list of particular client names, DCT versions, user agents, or classifications, by providing corresponding query parameters.

1. Filter the API calls by client names.

```
curl --location 'https://[Inser_DCT_Server]/v3/reporting/api-usage-report?
group_by=client_name%2Ckind&client_name=client1%2Cclient2' \
--header 'Authorization: <api_key>' \
--header 'Authorization: apk <api key>'
```

2. Filter the API calls by API classification.

```
curl --location 'https://[Inser_DCT_Server]/v3/reporting/api-usage-report?
group_by=client_name%2Ckind&api_metric_kind=automation' \
--header 'Authorization: <api_key>' \
--header 'Authorization: apk <api key>'
```



The API query parameters for this report are dynamic; depending on the number of records in the backend and the granularity of the response requested, API response can be too large to be handled by DCT. DCT can run out of memory and eventually crash if that is the case. To prevent this from happening, this report has the maximum limit set to 10,000 records in the API

response. Thus, it is strongly recommended to always ‘limit’ the scope of the usage response by filtering records on the basis of `start_date` and `end_date` query parameters, or on the basis of client names and/or user agents.

7.2 Central governance workflows

Data Control Tower delivers the management layer for all connected Delphix engines by virtue of its converged architecture. As such, DCT has the ability to simplify everyday administration of common engine admin tasks. This section will go over various how DCT exposes object relationships and reports on meaningful use patterns under insights.

- [Managing engines \(Continuous Data\)](#) (see page 162)
- [Managing dSources](#) (see page 170)
- [Managing VDBs](#) (see page 173)
- [Managing environments \(Continuous Data\)](#) (see page 176)
- [Managing bookmarks](#) (see page 178)
- [Replication management](#) (see page 179)

7.2.1 Managing engines (Continuous Data)

7.2.1.1 Infrastructure

DCT provides a near real-time list of all connected Continuous Data engines and lists them in an aggregate view. From the below screen, Delphix administrators can easily view and manage their engine connections.

The screenshot shows the 'Infrastructure' page in the Data Control Tower interface. The page title is 'Infrastructure' and it includes a 'Connect Engine' button in the top right corner. Below the title, there is a search bar and a 'Search (All)' dropdown. The main content is a table listing connected engines with columns for Status, Name/UUID, Type, Host, and Tags. The table contains three entries:

Status	Name/UUID ↑	Type	Host	Tags
Online	Test Engine 42366308-1323-71ff-e258-9c642ad8a2ae	Virtualization 9.0.0.0	lw01xyelphixlwmhy7.vm.cld.sr	Add Tags > View >
Online	amaresh-engine1 59454aa4-48b1-9455-7873-78833008418a	Unset 10.0.0.0	amaresh-engine1.dco1.delphix.com	Add Tags > View >
Online	bhavik-engine ec22420f-24e9-52fe-b8ac-fa3af4b7e7c	Virtualization 10.0.0.0	bhavik-engine.dlpdc.co	Add Tags > View >

Administrators can manage engine connects via the “Connect Engine” button on the top right corner. By clicking this button, the below dialogue will appear, asking for connection details.



DCT will access the engine as a registered user and, as detailed in the **Deployment** section, requires both a username and password as well as admin-level access to the engine.

Connect Engine



- Engine Details
- Authentication
- Security
- Tags
- Summary

Engine Details

Name
Test Engine

Hostname
uvo1xyellphikiwmhy7.vm.cid.sr

Choose Engine Type

- Virtualization
- Masking

7.2.1.2 Engine overview

Individual engine details can be seen and acted upon by clicking down on a particular engine detailed view. Once clicked, users will be sent to an "Overview" tab that provides relevant metadata related to the engine. From the Actions menu, you can also Deregister an engine.



Deregistering or deleting an engine can also be done via the API with:

```
curl --location --request DELETE 'https://<DCT Host Name>:<Port>/v3/management/engines/<ID or Name of the Engine to be deleted>' \
--header 'Authorization: apk <api key for authentication>'
```

The screenshot shows the 'Data Engines' overview page for a 'Test Engine'. The navigation bar includes 'Home', 'Data', 'Compliance', 'Insights', and 'Admin'. The breadcrumb trail is 'Data Engines > Test Engine'. The main content area features a storage usage chart showing 31.8% used storage (6.93GB) and 14.86GB remaining. Below the chart is a 'Details' table with columns for 'Type' and 'Version', showing 'Data Engine' and '14.0.0.0'. To the right is a 'Tags' section with an 'Edit' link. At the bottom right, a 'View All' button is visible. A red box highlights the 'Deregister' button in the 'Actions' menu.

7.2.1.2.1 Engine-connected environments

The "Environments" tab presents all environment connections to that particular engine.

Name ↑	Hosts IP (Port)	DBMS (Version)	Config	Engine ID	Tags
Oracle_Source	10.160.1.21 10.160.1.21	Unstructured Files mysql-plugin Oracle (19.3.0.0.0)	Standalone	3	Add Tags >
Oracle_Target	10.160.1.61 10.160.1.61	Unstructured Files mysql-plugin Oracle (19.3.0.0.0)	Standalone	3	Add Tags >
Postgres_Source	10.160.1.20 10.160.1.20	Unstructured Files postgres-vsdk	Standalone	3	Add Tags >
Postgres_Target	10.160.1.60 10.160.1.60	Unstructured Files postgres-vsdk	Standalone	3	Add Tags >
Sqsqlserver_Source	10.160.1.22	MSSql (14.0.2027.2)	Standalone	3	Add Tags >
Sqsqlserver_Target	10.160.1.62 10.160.1.62	Unstructured Files MSSql (14.0.2027.2)	Standalone	3	Add Tags >


7.2.1.2.2 Local dSources

The "dSources" tab presents all dSources associated with the selected engine. Clicking the "View" button will link the user directly to the associated dSource page.

Status	Name ↑	Type	Tags
UNKNOWN	AppFS_cust_master	Unstructured Files	Add Tags > View >
UNKNOWN	AppFS_master	Unstructured Files	Add Tags > View >
UNKNOWN	Mysql_master	mysql-plugin	Add Tags > View >
UNKNOWN	Oracle_master	Oracle	Add Tags > View >
UNKNOWN	Postgres_cust_master	postgres-vsdk	Add Tags > View >


7.2.1.2.3 Local VDBs

The "VDBs" tab presents all VDBs associated with the selected engine. Clicking the "View" button will link the user directly to the associated VDB page to take action.




7.2.1.2.4 Engine-based Operations access

Users are able to audit which users have access to this particular engine, what access group they belong to, and the associated permissions that each user has on this engine. Admins are able to click on the "View" button to access further details under the access control screen related to that specific user.



7.2.1.3 Infrastructure connection wizard

In the DCT interface, select the **Home** tab on the top navigation bar to see the Infrastructure Connections landing view, which hosts the list of connections. In DCT versions before 9.0.0, these connections were listed under the Environments section of the Data tab. Infrastructure Connections are the DCT equivalent of an environment. The **+ Infrastructure Connections** button will launch the wizard.



The following connections are available:

- UNIX environment
 - Standalone/Cluster
- Windows environment
 - Source/Target
 - Standalone/Cluster

7.2.1.3.1 Infrastructure connection wizard steps

1. **Connection Name**

Enter the **Connection Name** in the corresponding field and select the **Associated Engine** from the dropdown.

2. **Host & Server**

Select the **Host OS**, **Server Type**, and **Host Type** (applicable to Windows).

1 Unix

Create Infrastructure Connection
✕

- Connection Name
- Host & Server**
- Settings
- Summary

Host & Server

A Infrastructure Connection is a host or cluster with which Delphix will communicate.

Host OS

Unix/Linux

Windows

Host Type

Source

Target

Server Type

Standalone

Cluster

Cancel
Back
Next
Submit

2 Windows

For Windows/Target/Standalone settings, a Delphix Connector download link has been added. Unlike the engine, this link makes an API call to authenticate and download the connector exe file.

3. **Settings**

This step includes various connection setting options from basic items like the **Host/IP Address** and **SSH Port** to advanced items like **Discover SAP ASE, Provide my own JDK,** and **Set NFS**. It includes a **Validate** button to help confirm the environment user and prevents access to the next step if the credentials are not valid.

Create Infrastructure Connection
✕

- Host & Server
- Settings**
- Summary

A Connection is a host or cluster with which Delphix will communicate.

Toolkit Path

Java Development Kit

Provide my own JDK
Delphix will automatically provide a default JDK unless you elect to provide your own by checking 'Provide my own JDK'

Advanced Options

Discover SAP ASE

Login Settings

Username and Password

Password Vault

Kerberos

ASE DB Username

ASE DB Password

Set Network File System Addresses(NFS)
Setting NFS addresses is not normally needed. If the host has multiple IP addresses, these can be added for NFS traffic by selecting this option to add additional host addresses.

Set Delphix Session Protocol Options(DSP)
This option does not need to be selected unless you wish to configure Server/Client authentication for Delphix Session Protocol communication

Description

Cancel
Back
Next
Submit

A wide range of **Login Settings** are available for the OS user and to **Discover SAP ASE** (if applicable), like username/password, username/public key, Password Vault, or Kerberos.

- a. For vaults, HashiCorp and CyberArk vaults are supported.
- b. The Kerberos login option only shows up when the user selects a Kerberos enabled engine as the target for environment creation.

The screenshot shows the 'Create Infrastructure Connection' dialog box with the 'Settings' tab selected. The left sidebar contains a progress indicator with steps: Connection Name, Host & Server, Settings (selected), and Summary. The main content area is titled 'Settings' and includes a description: 'A Connection is a host or cluster with which Delphix will communicate.' Below this are input fields for 'Host/IP Address' and 'SSH Port' (with the value '22'). Under 'Login Settings', there are radio buttons for 'Username and Password' (selected), 'Username and Public Key', 'Password Vault', and 'Kerberos'. Below these are input fields for 'OS Username' and 'OS Password', a 'Validate' button, and a 'Toolkit Path' field. At the bottom, there is a 'Java Development Kit' section and a row of buttons: 'Cancel', 'Back', 'Next', and 'Submit'.

4. **Java Development Kit** (only applicable if selected in **Settings**)

Set the custom JDK path in the corresponding text field.

The screenshot shows the 'Create Infrastructure Connection' dialog box with the 'Java Development Kit' tab selected. The left sidebar contains a progress indicator with steps: Connection Name, Host & Server, Settings, Java Development Kit (selected), NFS Addresses, DSP Options, and Summary. The main content area is titled 'Java Development Kit' and includes the text: 'The current Java Development Kit that is installed: Default JDK'. Below this is a text input field labeled 'Java Development Kit (JDK) Path' containing the text 'path/to/jdk'. A note below the field reads: 'Provide the full (absolute) path to the root of the JDK.' At the bottom, there is a row of buttons: 'Cancel', 'Back', 'Next', and 'Submit'.

5. **NFS** (only applicable if selected in **Settings**)

Set NFS addresses in the corresponding text field (comma separated).

Create Infrastructure Connection
✕

- Connection Name
- Host & Server
- Settings
- Java Development Kit
- NFS Addresses**
- DSP Options
- Summary

NFS Addresses

Set Network File System (NFS) Addresses

NFS Addresses

Cancel Back Next Submit

6. **DSP** (only applicable if selected in **Settings**)
Set DSP configurations in this step.

Create Infrastructure Connection
✕

- Connection Name
- Host & Server
- Settings
- Java Development Kit
- NFS Addresses
- DSP Options**
- Summary

DSP Options

These fields have no effect unless client and server authorization for remote connections is enabled. These can be enabled in the Network Security section of the associated engine's Setup application dashboard.

DSP KeyStore Path

DSP KeyStore Password

DSP KeyStore Alias

DSP TrustStore Path

DSP TrustStore Password

Cancel Back Next Submit

7. **Summary**

Shows a comprehensive summary of the selected options in the previous configuration steps. Shows the type of login being used for both the OS user and SAP ASE (if applicable).

Create Infrastructure Connection
✕

- Connection Name
- Host & Server
- Settings
- Java Development Kit
- NFS Addresses
- DSP Options
- Summary

Summary

Review the configuration for this Infrastructure Connection

<p>Host</p> <p>Host OS Unix/Linux</p> <p>Server Type Standalone</p> <p>Connection Name My IC</p> <p>Associated Engine sj-2903.dcol1</p> <p>Host/IP Address unix-stand.dcol1.delphix.com</p> <p>NFS Addresses host1, host2.com, 1.12.21.32</p> <p>SSH Port 22</p> <p>Toolkit Path /work</p> <p>Java Development Kit (JDK) Path path/to/jdk</p> <p>DSP KeyStore Path DSP config1</p> <p>DSP KeyStore Alias DSP config2</p> <p>DSP TrustStore Path DSP config3</p> <p>Discover SAP ASE</p>	<p>User</p> <p>Login Type Password authentication used</p> <p>OS Username sybase</p>
---	---

Cancel Back Next Submit

7.2.2 Managing dSources

7.2.2.1 Managing dSources

DCT provides the ability to view, search, sort, and filter all dSources within a connected Delphix ecosystem. This page can be found under the **Data** section and is used to find and act upon all dSources, if they have the appropriate access.

Status	Name ↑	Type	Engine	Tags
UNKNOWN	AppFS_cust_master	Unstructured Files	Test Engine	Add Tags > View >
UNKNOWN	AppFS_master	Unstructured Files	Test Engine	Add Tags > View >
UNKNOWN	CDOMLOGR421FPDB1	Oracle	bhavik-engine	Add Tags > View >
UNKNOWN	CDOMLOG5813PDB1	Oracle	bhavik-engine	Add Tags > View >
RUNNING	CDOMSHRASEEPDB2:rhel-79-83ub-qar-74961-27a4593a.dool1.delphix.com	Oracle	amaresh-engine1	Add Tags > View >
UNKNOWN	DBOMSR8A1718	Oracle	bhavik-engine	Add Tags > View >
UNKNOWN	Mysqf_master	mysql-plugin	Test Engine	Add Tags > View >
UNKNOWN	Oracle_master	Oracle	Test Engine	Add Tags > View >
UNKNOWN	Postgres_cust_master	postgres-vsck	Test Engine	Add Tags > View >
UNKNOWN	Postgres_master	postgres-vsck	Test Engine	Add Tags > View >
UNKNOWN	Sutecrm_master	MSSql	Test Engine	Add Tags > View >
UNKNOWN	flaskapp	MSSql	Test Engine	Add Tags > View >

7.2.2.2 dSource overview

Individual dSource details can be viewed and acted upon by clicking down on a particular dSource's detailed view. Once clicked, users will be sent to an "overview" tab that provides relevant metadata related to the dSource.

Platform

Unstructured Files

Status

UNKNOWN

Details

Platform	Unstructured Files
Version	
Enabled	False
Size	5.20MB
Engine	Test Engine


Tags Edit

Environment Details

Environment Name	Postgres_Source
Type	Single Instance
OS	Linux

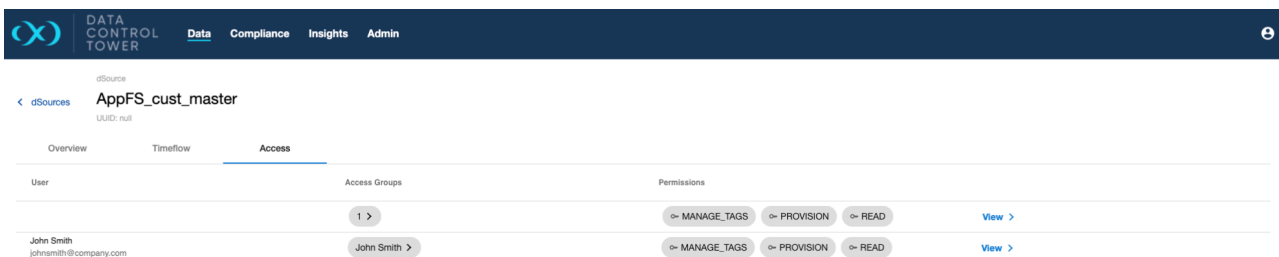
7.2.2.3 Timeflow visibility

Users are able to view snapshot information by tabbing over to the "timeflow" section, which lists all available snapshots via a vertical timeline. Users are able to modify snapshot retention periods by clicking on the ellipsis located to the right of the relevant snapshot.



7.2.2.4 Access auditing

Users are able to audit what other users have access to a particular dSource, what access group they belong to, and the associated permissions that each user has on that particular dSource.



7.2.2.5 dSource linking

DCT 10.0.0 introduces the feature to link a dSource using DCT APIs. With the new API parameters exposed from DCT, you can get the dSource linking defaults and use them to connect the dSource. The API parameters are available in the [API references](#)⁸⁰ article under:

- **BaseDataSourceLinkSourceParameters**
- **OracleDataSourceLinkSourceParameters**
- **OracleStagingPushDataSourceLinkSourceParameters**
- **AppDataDataSourceLinkSourceParameters**
- **ASEDataSourceLinkSourceParameters**
- **LinkDataSourceResponse**
- **LinkDataSourceDefaultRequest**
- **MSSQLDataSourceLinkSourceParameters**
- **MSSQLDataSourceStagingPushLinkSourceParameters**



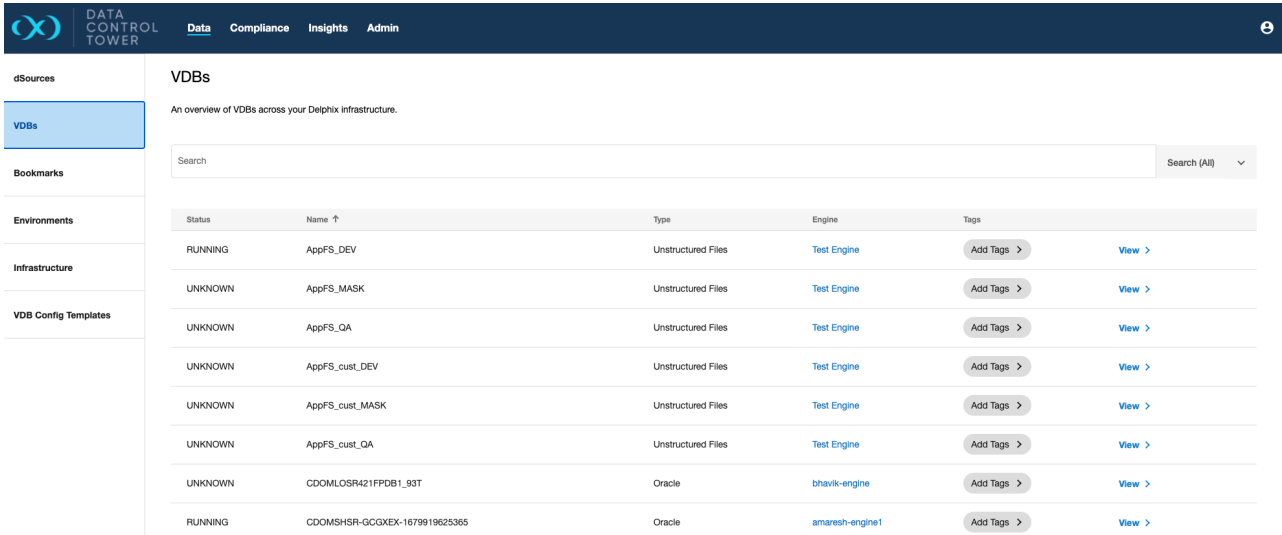
DCT supports APIs for linking all types of databases. DCT does not have a GUI flow for these APIs yet, which is why they are API only.

As a prerequisite, Continuous Data Engines should be registered in DCT and should have non-linked sources to be linked as a dSource.

⁸⁰ <https://dct.delphix.com/docs/latest/api-references>


7.2.3 Managing VDBs

DCT provides the ability to view, search, sort, and filter all VDBs within a connected Delphix ecosystem. This page can be found under the **Data** section and is used to find and act upon all VDB if they have the appropriate access.



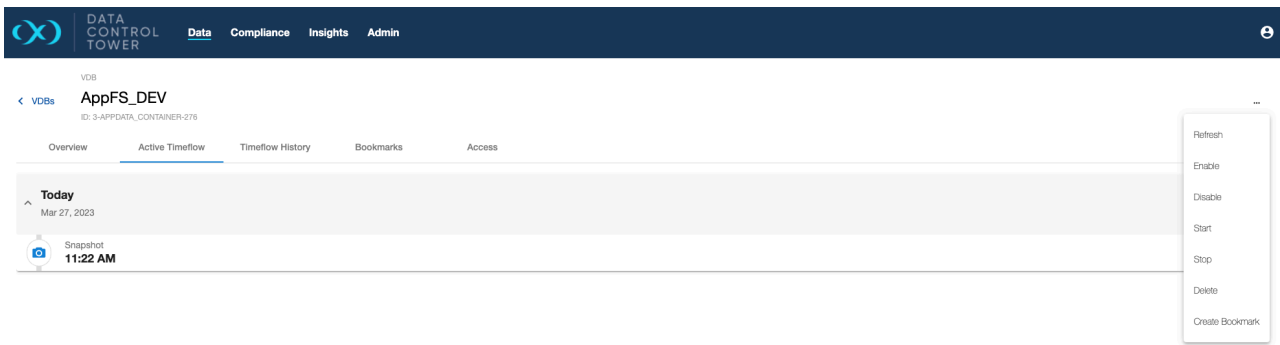
7.2.3.1 VDB overview

Individual VDB details can be seen and acted upon by clicking down on a particular VDB detailed view. Once clicked, users will be sent to an "overview" tab that provides relevant metadata related to the VDB.




7.2.3.2 VDB active timeline

Actionable snapshots are listed on the "active timeflow" tab - from this page, users can refresh, enable, disable, start, stop, delete, and create bookmarks on the VDB. Navigate to the Continuous Data workflows section, then VDB operations in the UI, and see Active timeline UI for more details.




7.2.3.3 VDB timeline history

A chronological history of all non-active timelines (commonly referred to as timeflows) is shown under the "Timeflow History" tab. From this page, developers can curate their QA or Development work by renaming timeflows to match their testing history. Developers also have the ability to access old timeflow data by making a particular timeflow "active". Navigate to the **Continuous Data workflows** section, then **VDB operations in the UI**, and see [Timeline history UI](#)⁸¹ for more details.



7.2.3.4 VDB bookmarks


A list of all bookmarks generated on the selected VDB can be found under the "Bookmarks" tab. This page provides a list of all bookmarks allowing for general organization and actions (developers can use bookmarks as a refresh or provision point from the API).



⁸¹ <https://portal.document360.io/docs/timeline-history>


7.2.3.5 VDB access

Users are able to audit which users have access to this particular VDB, what access group they belong to, and the associated permissions that each user has on that VDB. Admins are able to click on the "View" button to access further details under the access control screen related to that specific user.




7.2.3.6 VDB templates

Importing and removing imported VDB templates from connected engines is an available action from the **VDB Config Templates** page.




7.2.3.6.1 Import templates

To import a template, select the engines from the list in the import dialog.



7.2.3.6.2 Remove imported templates


To remove a template, select the engines from the list in the remove dialog.



7.2.4 Managing environments (Continuous Data)

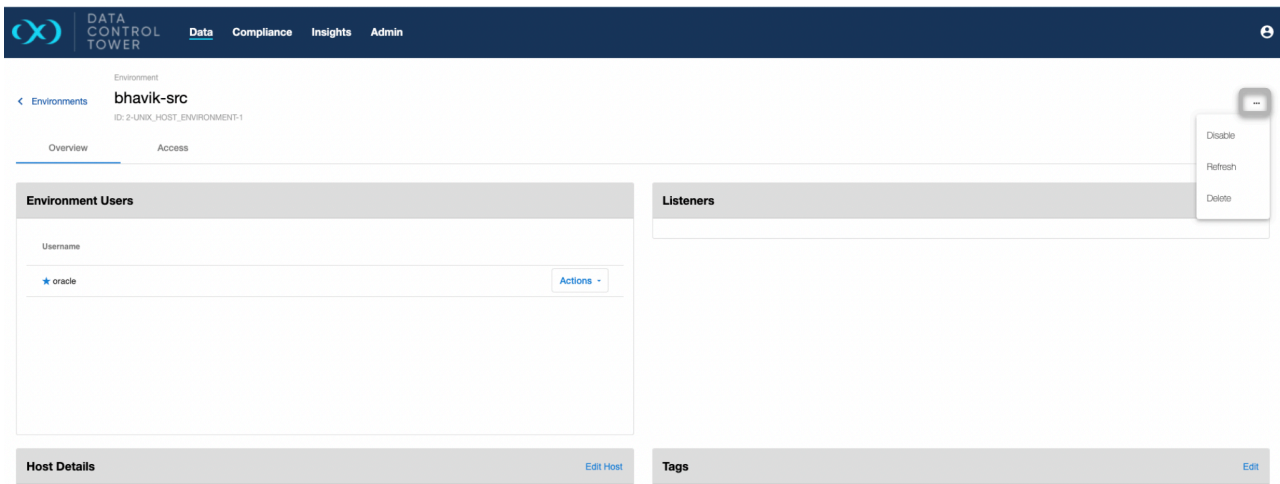
7.2.4.1 Global environments list

DCT provides the ability to view, search, sort, and filter all Continuous Data environments within a connected Delphix ecosystem. This page can be found under the **Data** section and is used to find and act upon all environment connections.



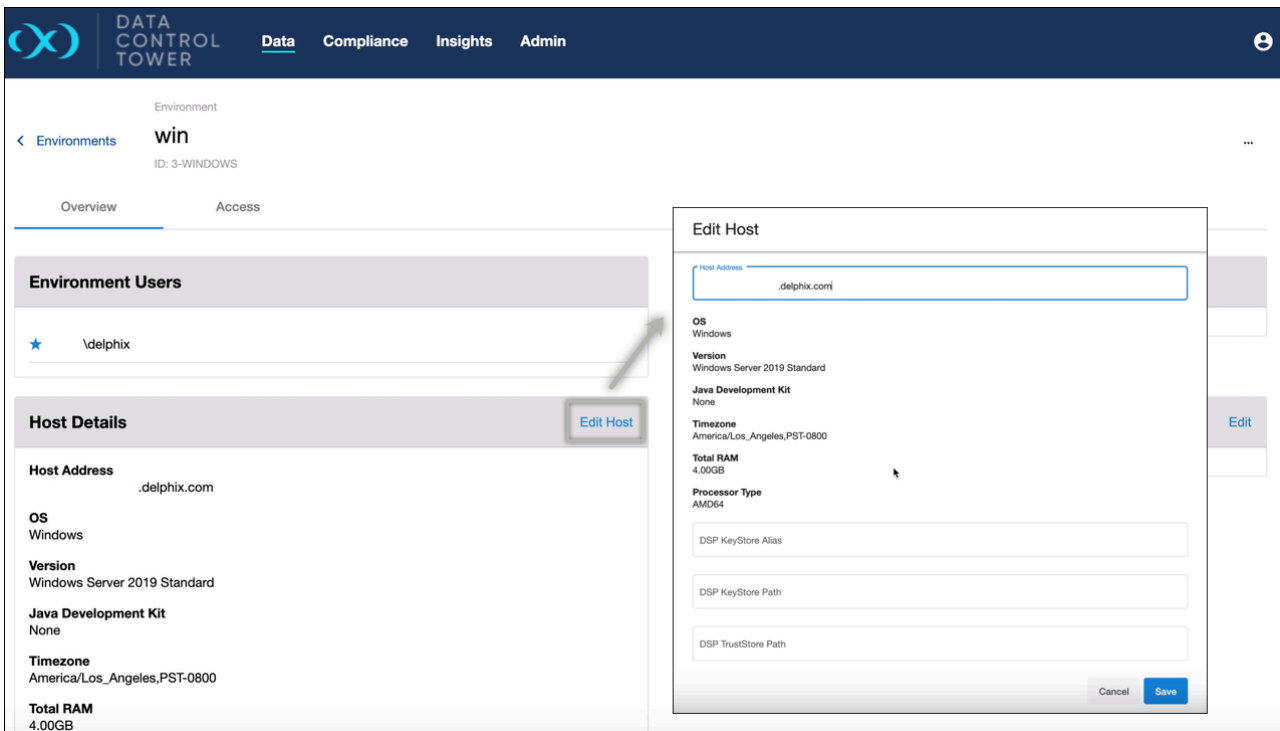
7.2.4.2 Manage environments

Selecting a standalone environment in the **Data** page shows an ellipsis in the top right corner. When the button is selected, the option to Enable/Disable, Refresh, or Delete the environment appears.



7.2.4.3 Edit host details

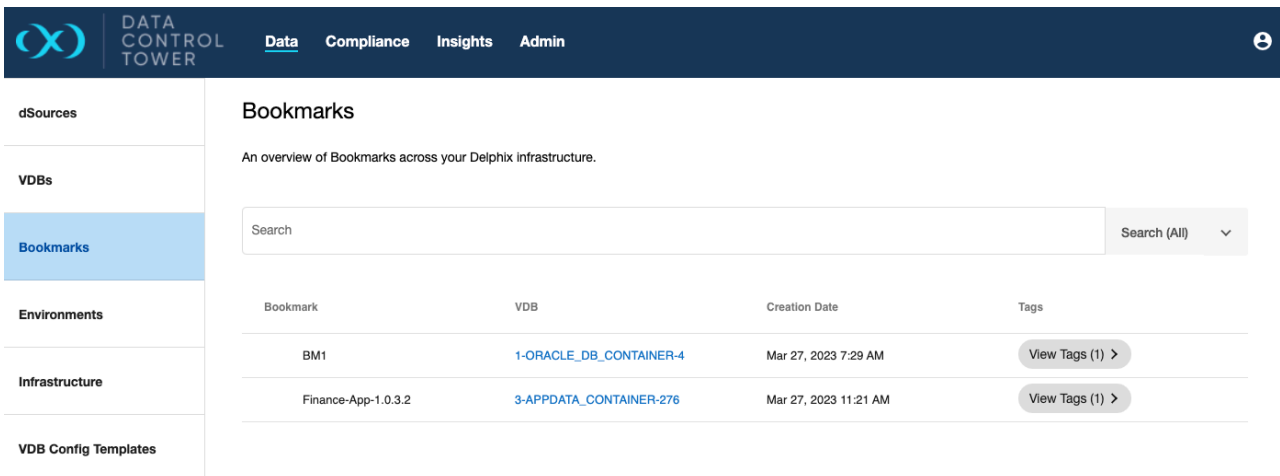
Selecting a standalone environment in the **Data** page shows an **Edit Host** option; it is not yet available for cluster environments. When the button is selected, the host details window will open, showing the input fields that can be edited. Select 'Save' to confirm the changes and close the window.



7.2.5 Managing bookmarks

7.2.5.1 Global Bookmarks List

DCT provides a near real-time list of all bookmarks across all VDBs and VDB-groups and their associated VDB(s). From the below screen, Delphix administrators can easily view and manage their bookmark estate.



Support has been added for the creation of bookmarks for dSource snapshots, similar to VDB snapshots. These properties have been added:

- `inherit_parent_tags` is a new property added in the request payload. This indicates whether this bookmark should inherit tags from the parent dataset.

- `dsource_ids` is a new property in the create bookmark response. This indicates the list of dSource IDs associated with this bookmark.

7.2.6 Replication management

7.2.6.1 Introduction

DCT 9.0.0 introduces the feature to differentiate between the replicated objects and original objects in case of a parent and replicated engines, both are registered with DCT.

7.2.6.2 Prerequisites

One parent and one replicated engine is required with few replicated dSources and VDBs.

7.2.6.3 User interface

- Convenient separation between the replicated objects(VDB/dSource/environments) and original objects.

The screenshot shows the 'VDBs' section of the Data Control Tower interface. The left sidebar contains a navigation menu with 'VDBs' selected. The main content area has a 'Provision VDB' button and a search bar. Below the search bar is a table with the following data:

Status	Name ↑	Type	Engine	Tags
RUNNING	CDOMLOSR421FPDB1_MCD	Oracle	amit-engine	Add Tags > View >
N/A	CDOMLOSR421FPDB1_MCD Replica	Oracle	replicated-engine	Add Tags > View >
RUNNING	CDOMLOSR421FPDB2_BAF	Oracle	amit-engine	Add Tags > View >
N/A	CDOMLOSR421FPDB2_BAF Replica	Oracle	replicated-engine	Add Tags > View >
RUNNING	DBOMSR8A1718_LJS	Oracle	amit-engine	Add Tags > View >
N/A	DBOMSR8A1718_LJS Replica	Oracle	replicated-engine	Add Tags > View >

At the bottom right of the table, there is a pagination control showing 'Items per page: 50' and '1 - 6 of 6'.

- Users can filter the replicated object using the advanced filter.

VDBs

An overview of VDBs across your Delphix infrastructure.

Search: (is_replica EQ true) Advanced Search

Advanced Search

Delphix supports advanced-based searching. See below to use the query builder.

SEARCH

Is Replica EQUAL TO True

Status	Name ↑	Type	Engine	Tags
N/A	CDOMLOSR421FPDB1_MCD Replica	Oracle	replicated-engine	Add Tags > View >
N/A	CDOMLOSR421FPDB2_BAF Replica	Oracle	replicated-engine	Add Tags > View >
N/A	DBOMSR8A1718_LJS Replica	Oracle	replicated-engine	Add Tags > View >

Items per page: 50 1 - 3 of 3

- All the actions on the replicated objects are disabled.

VDBs > CDOMLOSR421FPDB1_MCD
CDOMLOSR421FPDB1_MCD
ID: 2-ORACLE_DB_CONTAINER-37

Overview Active Timeflow Timeflow History Bookmarks Access

Platform: Oracle

Status: N/A

Actions: Refresh, Enable, Disable, Start, Stop, Delete, Create Bookmark

Details

Type	VDB
Platform	Oracle
Version	19.3.0.0.0
Size	0.00B
Engine	replicated-engine

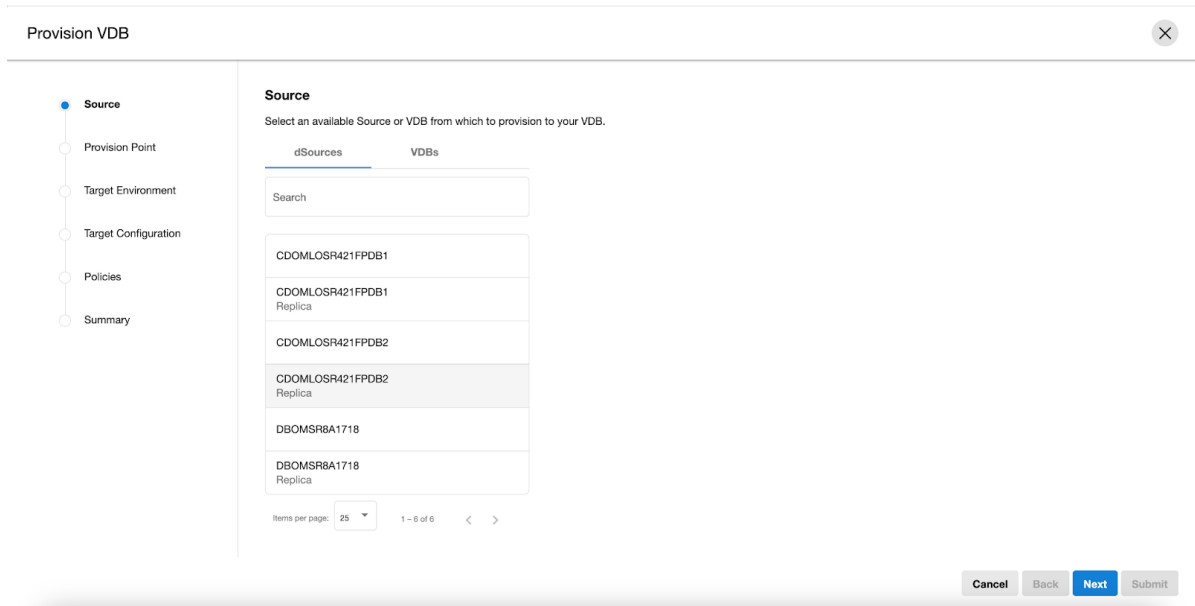
Tags Edit

No tags added

Environment Details

Environment Name	oracle-tgt
Type	Single Instance
OS	Linux

- The status of the replicated objects are N/A.
- On the VDB provisioning wizard, replicated objects(VDB/dSources) are marked as replicated.



7.2.6.4 API

Below are the list and search APIs, updated to return three additional fields (`is_replica` , `namespace_id` , `namespace_name`):

1. Sources

GET: `/sources`

GET: `/sources/{sourceId}`

POST: `/sources/search`

2. VDBs

GET: `/vdbs`

GET: `/vdbs/{vdbId}`

POST: `/vdbs/search`

3. dSources

GET: `/dsources`

GET: `/dsources/{dsourceId}`

POST: `/dsources/search`

4. CDBs

GET: `/cdbes`

GET: `/cdbes/{cdbId}`

POST: `/cdbes/search`

5. VCDBs

GET: /vcdb

GET: /vcdb/{vcdbId}

POST: /vcdb/search

6. Environments

GET: /environments

GET: /environments/{environmentId}

POST: /environments/search

7. Dataset-groups

GET: /groups

GET: /groups/{groupId}

POST: /groups/search

8. Timeflows

GET: /timeflows

GET: /timeflows/{timeflowId}

POST: /timeflows/search

9. Policies

GET: /virtualization-policies

GET: /virtualization-policies/{policyId}

POST: /virtualization-policies/search

10. Snapshots

GET: /snapshots

GET: /snapshots/{snapshotId}

POST: /snapshots/search

Below are the three fields added in response of these APIs:

1. `is_replica` : Boolean telling that this object id replicated or not.
2. `namespace_id` : This field will only come for replicated objects and is essentially the `namespace id` of the replicated object.
3. `namespace_name` : This field will only come for replicated objects and is essentially the `namespace name` of the replicated object.

7.2.6.5 Replication relationship

7.2.6.5.1 Overview

DCT 11.0 release adds a functionality to get the primary and replica objects for a given object in DCT. Whenever DCT identifies that a replica object has been added to the engine, it creates a job to fetch the primary object information from its source engine. In addition, when replica objects are updated with primary object information, the primary objects are also updated with their children (replicas).


This feature requires that both source and target engines are registered in DCT.

7.2.6.5.2 Prerequisites

A running instance of DCT, two instances of Delphix Engines, and some of the dSource and VDBs on one of those engines. A replication profile needs to be created from an engine that contains dSource and VDBs to replicate to another engine.

7.2.6.5.3 Replication details

A replication details box will appear on the VDB or dSource tab showing the **Source Engine** and **Source VDB/dSource** for replication.



7.2.6.5.4 API changes

New fields have been added in the following APIs:

- GET - /v3/dsources
- GET - /v3/dsources/{dsourceld}
- POST - /v3/dsources/search
- GET - /v3/vdbs

- GET - /v3/vdbs/{dsourceld}
- POST - /v3/vdbs/search

The new fields are as follows:

- primary_object_id
- primary_engine_id
- primary_engine_name
- replicas
 - replica_id
 - replica_engine_id
 - replica_engine_name
 - replica_namespace_id

7.2.6.5.4.1 VDB replica and primary object sample response:

```
curl --location 'http://localhost:8080/v3/vdbs' \
--header 'Authorization: apk {{authToken}}'

{
  "items": [
    {
      "id": "1-ORACLE_DB_CONTAINER-11",
      "database_type": "Oracle",
      "name": "VCDO_IF0",
      "namespace_id": "1-NAMESPACE-2",
      "namespace_name": "ip-10-110-221-77-1",
      "is_replica": true,
      "is_locked": false,
      "database_version": "19.3.0.0.0",
      "size": 0,
      "storage_size": 2778624,
      "engine_id": "1",
      "masked": false,
      "content_type": "PDB",
      "parent_timeflow_timestamp": "2023-10-16T10:11:02Z",
      "parent_timeflow_timezone": "America/New_York,EDT-0400",
      "environment_id": "1-UNIX_HOST_ENVIRONMENT-3",
      "ip_address": "10.43.89.210",
      "fqdn": "ora-src.dcol1.delphix.com",
      "parent_id": "1-ORACLE_DB_CONTAINER-8",
      "parent_dsourceld": "1-ORACLE_DB_CONTAINER-8",
      "group_name": "Untitled",
      "engine_name": "e1",
      "cdb_id": "1-ORACLE_SINGLE_CONFIG-24",
      "creation_date": "2023-10-16T10:38:39.05Z",
      "hooks": {
        "pre_refresh": [],

```



```

    "post_refresh": [],
    "pre_self_refresh": [],
    "post_self_refresh": [],
    "pre_rollback": [],
    "post_rollback": [],
    "configure_clone": [],
    "pre_snapshot": [],
    "post_snapshot": [],
    "pre_start": [],
    "post_start": [],
    "pre_stop": [],
    "post_stop": []
  },
  "config_params": {
    "_cdb_disable_pdb_limit": "TRUE",
    "audit_file_dest": "'/u01/app/oracle/admin/CDOMLOSR421F/adump'",
    "audit_trail": "'DB'",
    "compatible": "'19.0.0'",
    "diagnostic_dest": "'/u01/app/oracle'",
    "dispatchers": "'(PROTOCOL=TCP) (SERVICE=CDOMLOSR421FXDB)'",
    "enable_pluggable_database": "TRUE",
    "log_archive_format": "'%t_%s_%r.dbf'",
    "max_pdfs": "4098",
    "memory_max_target": "1342177280",
    "memory_target": "1342177280",
    "nls_language": "'AMERICAN'",
    "nls_territory": "'AMERICA'",
    "open_cursors": "300",
    "processes": "300",
    "remote_login_passwordfile": "'EXCLUSIVE'"
  },
  "mount_point": "/mnt/provision",
  "current_timeflow_id": "1-ORACLE_TIMEFLOW-11",
  "vdb_restart": false,
  "is_appdata": false,
  "primary_object_id": "1-ORACLE_DB_CONTAINER-6",
  "primary_engine_id": "1",
  "primary_engine_name": "e1"
},
{
  "id": "1-ORACLE_DB_CONTAINER-6",
  "database_type": "Oracle",
  "name": "VCDO_IF0",
  "is_replica": false,
  "is_locked": false,
  "database_version": "19.3.0.0.0",
  "size": 794755072,
  "storage_size": 32859648,
  "engine_id": "1",
  "status": "RUNNING",
  "masked": false,
  "content_type": "PDB",

```

```

"parent_timeflow_timestamp": "2023-10-16T10:11:02Z",
"parent_timeflow_timezone": "America/New_York,EDT-0400",
"environment_id": "1-UNIX_HOST_ENVIRONMENT-1",
"ip_address": "10.43.89.210",
"fqdn": "ora-src.dcol1.delphix.com",
"parent_id": "1-ORACLE_DB_CONTAINER-2",
"parent_dsource_id": "1-ORACLE_DB_CONTAINER-2",
"group_name": "Untitled",
"engine_name": "e1",
"cdb_id": "1-ORACLE_SINGLE_CONFIG-2",
"creation_date": "2023-10-16T10:38:39.05Z",
"hooks": {
  "pre_refresh": [],
  "post_refresh": [],
  "pre_self_refresh": [],
  "post_self_refresh": [],
  "pre_rollback": [],
  "post_rollback": [],
  "configure_clone": [],
  "pre_snapshot": [],
  "post_snapshot": [],
  "pre_start": [],
  "post_start": [],
  "pre_stop": [],
  "post_stop": []
},
"config_params": {
  "_cdb_disable_pdb_limit": "TRUE",
  "audit_file_dest": "'/u01/app/oracle/admin/CDOMLOSR421F/adump'",
  "audit_trail": "'DB'",
  "compatible": "'19.0.0'",
  "diagnostic_dest": "'/u01/app/oracle'",
  "dispatchers": "'(PROTOCOL=TCP) (SERVICE=CDOMLOSR421FXDB)'",
  "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'"
},
"mount_point": "/mnt/provision",
"current_timeflow_id": "1-ORACLE_TIMEFLOW-6",
"vdb_restart": false,
"is_appdata": false,
"replicas": [
  {
    "replica_id": "1-ORACLE_DB_CONTAINER-11",
    "replica_engine_id": "1",

```

```

                "replica_engine_name": "e1",
                "replica_namespace_id": "1-NAMESPACE-2"
            }
        ]
    },
    "response_metadata": {
        "total": 2
    }
}

```

7.2.6.5.4.2 Sample response for dSource change:

```

curl --location 'http://localhost:8080/v3/dsources/search' \
--header 'Content-Type: application/json' \
--header 'Authorization: apk {{authToken}}' \
--data '{
  "filter_expression" : "engine_id EQ \'\'1\'\'"'
}'

{
  "items": [
    {
      "id": "1-ORACLE_DB_CONTAINER-10",
      "database_type": "Oracle",
      "name": "CDOMSHSR6706PDB2-NO-CHILD",
      "namespace_id": "1-NAMESPACE-2",
      "namespace_name": "ip-10-110-221-77-1",
      "is_replica": true,
      "database_version": "19.3.0.0.0",
      "content_type": "PDB",
      "data_uuid": "46f1613b223b1dd5364bdc3ecfd2755d",
      "storage_size": 180355584,
      "creation_date": "2023-10-16T10:09:00.567Z",
      "group_name": "Untitled",
      "engine_id": "1",
      "source_id": "1-ORACLE_PDB_CONFIG-20",
      "engine_name": "e1",
      "cdb_id": "1-ORACLE_SINGLE_CONFIG-23",
      "current_timeflow_id": "1-ORACLE_TIMEFLOW-7",
      "is_appdata": false,
      "primary_object_id": "1-ORACLE_DB_CONTAINER-4",
      "primary_engine_id": "1",
      "primary_engine_name": "e1"
    },
    {
      "id": "1-ORACLE_DB_CONTAINER-12",
      "database_type": "Oracle",

```

```

"name": "CDOMLOSR421FPDB2",
"is_replica": false,
"database_version": "19.3.0.0.0",
"content_type": "PDB",
"data_uuid": "837bbc258a479ff86dbce5656beeba2a",
"storage_size": 236433920,
"creation_date": "2023-10-17T10:19:19.496Z",
"group_name": "Untitled",
"enabled": true,
"engine_id": "1",
"source_id": "1-ORACLE_PDB_CONFIG-7",
"status": "RUNNING",
"engine_name": "e1",
"cdb_id": "1-ORACLE_SINGLE_CONFIG-2",
"current_timeflow_id": "1-ORACLE_TIMEFLOW-12",
"is_appdata": false
},
{
  "id": "1-ORACLE_DB_CONTAINER-2",
  "database_type": "Oracle",
  "name": "CDOMLOSR421FPDB1",
  "is_replica": false,
  "database_version": "19.3.0.0.0",
  "content_type": "PDB",
  "data_uuid": "834a794e51ec096059922dc06625fb66",
  "storage_size": 245758976,
  "creation_date": "2023-10-16T10:08:35.425Z",
  "group_name": "Untitled",
  "enabled": true,
  "engine_id": "1",
  "source_id": "1-ORACLE_PDB_CONFIG-6",
  "status": "RUNNING",
  "engine_name": "e1",
  "cdb_id": "1-ORACLE_SINGLE_CONFIG-2",
  "current_timeflow_id": "1-ORACLE_TIMEFLOW-2",
  "is_appdata": false,
  "replicas": [
    {
      "replica_id": "1-ORACLE_DB_CONTAINER-8",
      "replica_engine_id": "1",
      "replica_engine_name": "e1",
      "replica_namespace_id": "1-NAMESPACE-2"
    }
  ]
},
{
  "id": "1-ORACLE_DB_CONTAINER-21",
  "database_type": "Oracle",
  "name": "DSOURCE-WITH-CUSTOM-POLICY",
  "is_replica": false,
  "database_version": "19.3.0.0.0",
  "content_type": "PDB",

```

```

    "data_uuid": "dd4216e99ee17e8a7afa97cce012aa42",
    "storage_size": 181809664,
    "creation_date": "2023-10-18T13:09:30.299Z",
    "group_name": "Untitled",
    "enabled": true,
    "engine_id": "1",
    "source_id": "1-ORACLE_PDB_CONFIG-11",
    "status": "RUNNING",
    "engine_name": "e1",
    "cdb_id": "1-ORACLE_SINGLE_CONFIG-3",
    "current_timeflow_id": "1-ORACLE_TIMEFLOW-21",
    "is_appdata": false
  },
  {
    "id": "1-ORACLE_DB_CONTAINER-4",
    "database_type": "Oracle",
    "name": "CDOMSHSR6706PDB2-NO-CHILD",
    "is_replica": false,
    "database_version": "19.3.0.0.0",
    "content_type": "PDB",
    "data_uuid": "46f1613b223b1dd5364bdc3ecfd2755d",
    "storage_size": 190537728,
    "creation_date": "2023-10-16T10:09:00.567Z",
    "group_name": "Untitled",
    "enabled": true,
    "engine_id": "1",
    "source_id": "1-ORACLE_PDB_CONFIG-10",
    "status": "RUNNING",
    "engine_name": "e1",
    "cdb_id": "1-ORACLE_SINGLE_CONFIG-3",
    "current_timeflow_id": "1-ORACLE_TIMEFLOW-4",
    "is_appdata": false,
    "replicas": [
      {
        "replica_id": "1-ORACLE_DB_CONTAINER-10",
        "replica_engine_id": "1",
        "replica_engine_name": "e1",
        "replica_namespace_id": "1-NAMESPACE-2"
      }
    ]
  },
  {
    "id": "1-ORACLE_DB_CONTAINER-8",
    "database_type": "Oracle",
    "name": "CDOMLOSR421FPDB1",
    "namespace_id": "1-NAMESPACE-2",
    "namespace_name": "ip-10-110-221-77-1",
    "is_replica": true,
    "database_version": "19.3.0.0.0",
    "content_type": "PDB",
    "data_uuid": "834a794e51ec096059922dc06625fb66",
    "storage_size": 226567680,

```

```

    "creation_date": "2023-10-16T10:08:35.425Z",
    "group_name": "Untitled",
    "engine_id": "1",
    "source_id": "1-ORACLE_PDB_CONFIG-22",
    "engine_name": "e1",
    "cdb_id": "1-ORACLE_SINGLE_CONFIG-24",
    "current_timeflow_id": "1-ORACLE_TIMEFLOW-9",
    "is_appdata": false,
    "primary_object_id": "1-ORACLE_DB_CONTAINER-2",
    "primary_engine_id": "1",
    "primary_engine_name": "e1"
  }
],
"response_metadata": {
  "total": 6
}
}

```

7.3 Insight reports

7.3.1 Central governance insights

DCT provides global reporting of real-time statuses. This section will break down all of the reports in the **Insights** section of DCT.



All insight dashboards can be exported to CSV or JSON format.

7.3.2 VDB Inventory

The VDB Inventory report provides users with a comprehensive list of all the Virtual Databases (VDBs) created in the DCT platform and their identification metadata.

7.3.3 dSource Inventory

The dSource Inventory report provides users with a comprehensive list of all the dSources created in the DCT platform and their identification metadata.

7.3.4 Source Ingestion Metrics

The Data Source Ingestion Metrics dashboard is designed to help users find their virtualization source ingestion metrics, which are often required for contract renewal purposes. The dashboard contains the data sources and informs the user of the total size of that ingestion source.

7.3.5 Compliance Job Executions report

The Compliance Job Executions Report dashboard offers a comprehensive overview of the compliance jobs executed within DCT, by presenting key metrics that enable stakeholders to assess the efficiency and effectiveness of their data governance efforts.

7.3.6 Block Storage

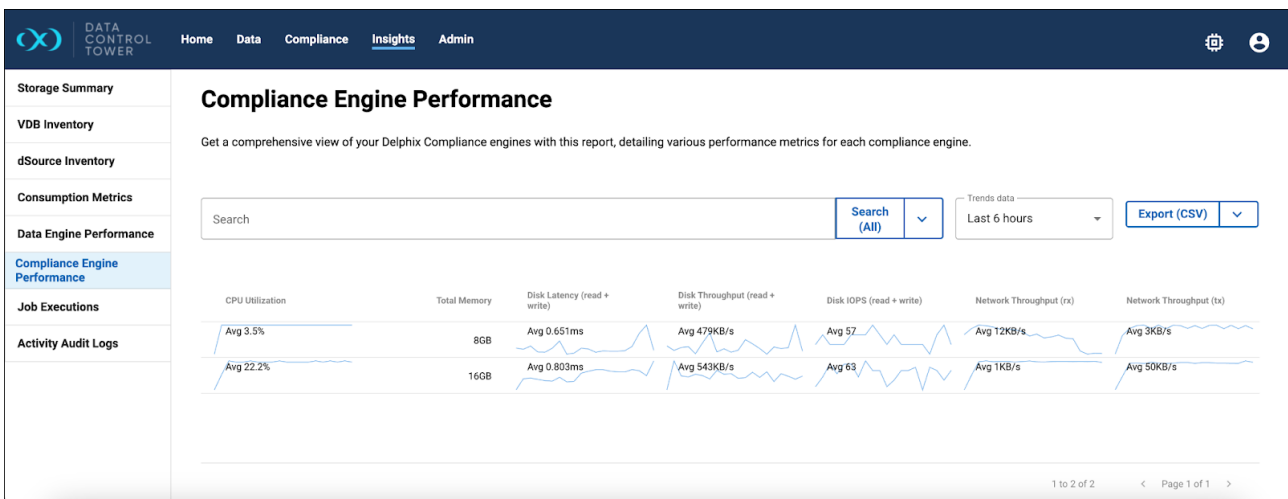
This report provides users with a comprehensive view of storage usage across different engines. With this report, users can easily identify the engines that are utilizing the most storage and take necessary action to optimize storage usage.

7.3.7 Activity Audit Log summary

The Activity Audit Log Summary provides a high-level audit log summary capturing the utilization of DCT by displaying user activity and the historical count of actions executed within the platform. This concise report enables stakeholders to quickly identify trends, monitor user engagement, and assess the overall effectiveness of data governance processes.


7.3.8 Compliance Engine Performance report

The Engine Performance dashboard provides users with a comprehensive view of performance information for Delphix Continuous Compliance Engines. The performance metrics include system CPU utilization, system memory usage, latency and throughput details for Disk protocols, and network throughput.




7.3.9 Data Engine Performance report

The Engine Performance dashboard provides users with a comprehensive view of performance information for Delphix Continuous Compliance Engines. The performance metrics include system CPU utilization, system memory usage, latency, IOPS, and network throughput details for Disk, NFS, and iSCSI protocols.



By default, a fix set of data columns are shown. Users have the flexibility to add or remove columns based on their preferences.



7.4 Tabular customization support for DCT

7.4.1 Introduction

The Data Control Tower (DCT) integration of AgGrid helps provide a more dynamic and customizable data interaction experience that enhances the way you interact with data tables in DCT.

7.4.1.1 Key enhancements

1. **Seamless integration with AgGrid:**

- **Dynamic data tables:** AgGrid, a leading data grid tool, is being progressively integrated across DCT. This transition promises a smoother, faster, and more responsive interaction with data.

2. **User-centric customization:**

- **Control at your fingertips:** AgGrid enables you to dynamically show or hide columns in tables, resize them as per your requirement, and rearrange them to suit your analysis process.


8 Continuous Data workflows

8.1 DevOps TDM

DCT delivers all of the Continuous Data and developer operations necessary to power DevOps and Test Data Management use-cases. This includes a suite of APIs to drive automation.



Using the above APIs, DCT can seamlessly integrate Delphix data into DevOps pipelines by providing a single point of integration for a broad Delphix deployment.



8.2 Developer experience

In addition to automation use cases, DCT provides the APIs and UI to power developer access to Delphix data and common Delphix operations. This section will detail all of the major capabilities that make up this revamped Delphix developer experience.

8.3 Self-service vs. DCT developer experience

Data Control Tower now provides a central experience for developers. Whether a developer prefers to leverage Delphix via API, integration, or UI, DCT delivers the ability to quickly access data from any connected Delphix engine, and the common capabilities to drive application development and testing.

Previously, Delphix offered a local add-on application called Self-Service (or Jet Stream) that was attached to applicable data engines. Self-Service provided an interface to access pre-provisioned datasets encapsulated in "Self-Service containers", which would be made available by admin configuration.

Data Control Tower has taken the most common operations and use-cases, and has made this experience accessible to developers via API, integration, and UI. This article will describe the key use-case and operational overlap, as well as the differences between the local engine Self-Service experience and DCT's developer experience.

8.3.1 Key similarities

1. Developer access to Delphix Data

The DCT developer experience is geared toward driving access to data, with all of the same time-based operations to enable application development and testing. Operations (accessible via the API, integration, or UI) include refresh, rewind, start/stop, enable/disable, bookmark, bookmark share, and timeflow visibility/access.

2. Developer timeflow history

A common UI benefit in Self-Service is the ability to visualize past timeflows (see [Timeline history](#) (see [page 217](#)) for more detail), which acts like a testing record. Every time a developer runs a test and rewinds/refreshes, that past test results are stored in Delphix as a timeflow. DCT has both API and UI instrumentation to make the visualization and curation of timeflows incredibly simple.

3. Data-as-Code

Developers can use DCT bookmarks to reference a point in time on a VDB (or group of VDBs) with a developer-set retention period and human-readable name. This is valuable for development teams as they evolve application code. Whenever a code change necessitates a new database schema, a developer can bookmark a VDB that is formatted to work with that particular code branch. This empowers development teams to always have access to a viable test data set for any and code branches of an application.

8.3.2 Key differences

1. DCT delivers a central interface powered by its converged architecture

This means that developers have a single location to log into in order to access and manipulate their virtual data sets.

2. User experience

The DCT developer experience UI has completely been reworked to make developer access to Delphix data easy and intuitive. This experience shows itself in three UI tabs, **Active Timeline**, **Timeline History**, and **Bookmarks**, that are located in each VDB's detail menu. This experience is meant to be used by all Delphix users (admins and developers, especially) and will be tailored to the individual based on the DCT Access Control system.

3. No template/container model

Previously, engine administrators needed to create templates encapsulating one or more related VDBs and provision new VDBs into a developer-accessible container. This model required manual administration that created bottlenecks for data access, which was especially prohibitive for automation use-cases. The benefit of this model was two-fold: **first**, containers represented a miniature sandbox for developers (using a Self-Service user role) and **second**, bulk operations could be performed on all container-grouped VDBs while maintaining referential synchronicity, a valuable attribute for integration testing.

4. DCT Access Control replaces the developer sandbox enabled by Self-Service containers

Developers simply log into DCT and can view and act upon data that they are entitled to access with

operations tightly bounded by their defined role. DCT's Access Control system has the ability to automate both user membership of access groups and entitlement access via attribute-defined scoped roles. In addition, roles can be customized in DCT such that granular permissions can be extended and restricted down to both access group and user levels.

5. **DCT VDB Groups replace the Self-Service container grouping mechanism**

Currently only available via API, VDB groups enable the association of one or more VDBs for bulk operations while maintaining referential synchronicity.

6. **Time operations consolidation**

The developer experience UI consolidates the many time-based operations across Continuous Data and Self-Service (e.g. refresh, rewind, rollback, restore, reset, etc.) into a single operation; **refresh**. From the DCT UI, clicking refresh will take users to a contextualized screen that simplifies time operations by focusing on what timeline (and what time) the user would like to align to (parent, self, or relative).

7. **No "branching"**

Branching in Self-Service introduced the notion of task-specific timelines, each with its own associated sets of timeflows. This was a concept that was heavily tied to the "template/container" model and is obviated by the DCT Access Control system that can enable gated provisioning access to a developer. If a new timeline is needed for a separate task, you can provision a new VDB.



DCT has a Delphix-supported integration with ServiceNow, which is commonly used as a developer resource-request tool. Users can build custom developer-centric workflows with any operation currently instrumented through the DCT API layer.

8.4 Creating and managing bookmarks

8.4.1 Create a Bookmark

Bookmarks are a critical developer tool that enables the creation of a namable time reference to a snapshot of a VDB or VDB group. Bookmarks for single VDBs can be created from the DCT UI by selecting a VDB and expanding into its detailed view. From the **Active Timeline** view, users can select the ellipsis in the top right corner and "Create Bookmark".

VDB

< VDBs **AppFS_DEV**
ID: 3-APPDATA_CONTAINER-276

Overview **Active Timeflow** Timeflow History Bookmarks Access

Today
Mar 27, 2023

Snapshot
11:22 AM

2 months ago
Jan 20, 2023

Refresh
Enable
Disable
Start
Stop
Delete
Create Bookmark

Selecting the "Create Bookmark" button will open a window that enables bookmark naming, setting the custom retention period for that bookmark, and assigning any relevant tags. Creating bookmarks this way will initiate a **new** snapshot operation that will then be associated with that bookmark.

Bookmarks relating to a specific VDB can be found under the bookmarks tab in a VDBs details page. This provides a curated list of actionable snapshots that represent anything from a relevant test result to a transformed set of schema that can be associated with a specific branch of code.

VDB

< VDBs **AppFS_DEV**
ID: 3-APPDATA_CONTAINER-276

Overview Active Timeflow Timeflow History **Bookmarks** Access


Bookmark	Creation Date	Tags	Actions
Finance-App-1.0.3.2	Mar 27, 2023 11:21 AM	View Tags (1) >	:

8.4.1.1 Create a bookmark from an existing snapshot


Starting in version 7.0, the UI has an option to create bookmarks from existing snapshots.

On the VDB detail page, under the **Active Timeline** tab for each snapshot, a **Create Bookmark** action is available. This opens a dialogue that shows a list of inputs for the user to select from, to create a bookmark.


Once the user clicks **Create Bookmark** in the dialogue, the bookmark will be created for that particular snapshot (if all the mandatory fields are completed), else errors will be shown.



Create a bookmark at the **current** data and time.



Create a bookmark based on the **selected** date and time.



Create a bookmark based on **SCN Number**.

8.4.2 Bookmark API Documentation

Some advanced bookmark operations are only available via API at present, formal documentation can be found via DCT's swagger docs or the Developer resources section. This portion of the bookmarks documentation will discuss examples of advanced use-cases.


8.4.3 Create a Bookmark at the current time for multiple VDBs

DCT (version 6.0 and above) can create bookmarks from existing snapshots. This is particularly useful for users looking to migrate Self-Service bookmarks to DCT or any developer looking to retroactively create a bookmark reference.

In cases such as integration testing, bundling multiple VDBs together to represent a complete set of data that a complex application would run on is helpful. This API example shows how a single bookmark reference can be created off of multiple VDBs, to provide a provision point for new testing sets or the creation of a VDB Group that can be used to maintain referential synchronicity from that bookmark point.

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/bookmarks' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
```

```
"name": "MyBookmark1",
"vdb_ids": [
  "1-ORACLE_DB_CONTAINER-2",
  "2-ORACLE_DB_CONTAINER-2"
]
}'
```

 These API calls will return a DCT job to track the creation process. This job ID can then be used to poll the status via the jobs API. Example response:

```
{
  "bookmark": {
    "id": "9e8c7223f1af4694a19ac2c2f7696eda",
    "name": "MyBookmark1",
    "creation_date": "2023-03-27T20:56:13.916857Z",
    "vdb_ids": [
      "1-ORACLE_DB_CONTAINER-2",
      "2-ORACLE_DB_CONTAINER-2"
    ],
    "retention": 30,
    "expiration": "2023-04-26"
  },
  "job": {
    "id": "8fe825f5635d45299915c3cb88a17623",
    "status": "PENDING",
    "type": "BOOKMARK_CREATE",
    "target_id": "9e8c7223f1af4694a19ac2c2f7696eda",
    "start_time": "2023-03-27T20:56:14.363549Z"
  }
}
```

8.4.4 Creating a bookmark from a chosen timepoint

Now bookmarks can be created at a chosen snapshot timepoint using new properties in the request payload. Some properties are mutually exclusive, so be sure to use a valid request payload.

Here are some examples of a valid request payload:

```
{
  "name": "my-bookmark-123",
  "vdb_ids": [
    "vdb-123"
  ],
  "location": "112233",
  "expiration": "2021-07-04",
  "retain_forever": false,
```

```

"tags": [
  {
    "key": "key-1",
    "value": "value-1"
  },
  {
    "key": "key-2",
    "value": "value-2"
  }
],
"make_current_account_owner": true
}

```

```

{
  "name": "my-bookmark-123",
  "timeflow_ids": [
    "timeflow-1"
  ],
  "location": "112233",
  "expiration": "2021-07-04",
  "retain_forever": false,
  "tags": [
    {
      "key": "key-1",
      "value": "value-1"
    },
    {
      "key": "key-2",
      "value": "value-2"
    }
  ],
  "make_current_account_owner": true
}

```

```

{
  "name": "my-bookmark-123",
  "timeflow_ids": [
    "timeflow-1"
  ],
  "timestamp": "2021-05-01T08:51:34.148000+00:00",
  "expiration": "2021-07-04",
  "retain_forever": false,
  "tags": [
    {
      "key": "key-1",
      "value": "value-1"
    },
    {
      "key": "key-2",
      "value": "value-2"
    }
  ]
}

```

```

    }
  ],
  "make_current_account_owner": true
}

```

Properties

- timestamp
- location
- timestamp_in_database_timezone
- timeflow_ids

Rules

- timestamp , location and timestamp_in_database_timezone are mutually exclusive.
- If any of the properties from timestamp , location and timestamp_in_database_timezone are provided then:
 - At least one property for vdb_ids or timeflow_ids must be provided.
 - The snapshot_ids property is not allowed to be set.

8.5 VDB operations



The developer experience will continue to see investment and additional capabilities over the next few releases.

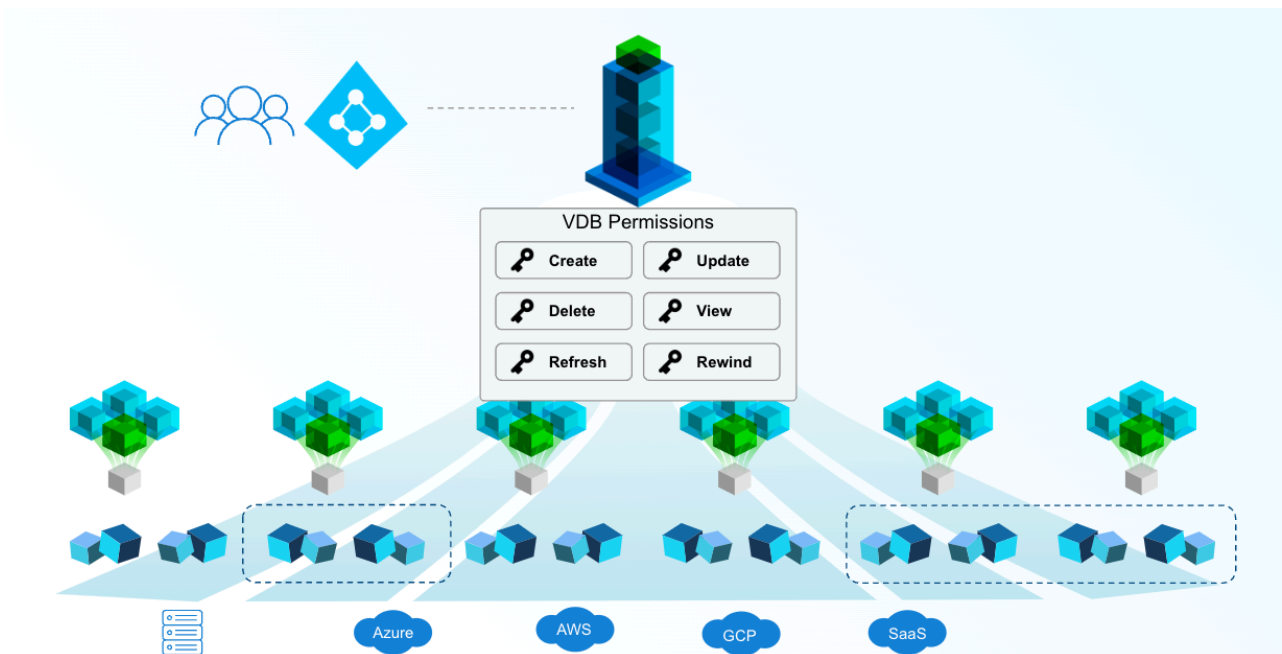
The VDB operations UI serves as an actionable command center for admins and developers. With this UI, users can migrate from using the local engine UI to leveraging DCT to do their daily VDB-related work. This encompasses both continuous data as well as any developers leveraging Delphix Self Service.

The core benefit of this UI experience is the breadth of access coupled with DCT's access control system. Using both together, a user can access and act upon any data on any connected engine within the boundaries of the entitlement and permissions set by the admin-driven access control system.

Status	Name ↑	Type	Engine	Compliance	Last Refresh	Tags
UNKNOWN	ChildVDB	Oracle	SJ	Not Masked	Sep 19, 2023 1:04 PM GMT+5:30	Add Tags > View >
UNKNOWN	DBOMSR91F858dSource2_VLN	Oracle	Engine	Not Masked		Add Tags > View >
UNKNOWN	EmptyVDB_90NC2M95	Unstructured Files	Engine	Not Masked		Add Tags > View >
UNKNOWN	JK_CD0MSHSR364EPDB31109_ZKD	Oracle	SJ	Not Masked		Add Tags > View >
UNKNOWN	VAppdata dSource_CEdrbQhH	Second_Plugin	DJ	Not Masked		Add Tags > View >
UNKNOWN	VAppdata dSource_GpqlRueN	Second_Plugin	DJ	Not Masked		Add Tags > View >
UNKNOWN	VAppdata dSource_IDnCsFtt	Second_Plugin	DJ	Not Masked		Add Tags > View >

To access the VDB operations UI, users only need to log into DCT and select the detailed view of any particular VDB.

Users will only be able to see VDBs if they have been granted access via the Access Control system.



From there, users can perform common operations such as refresh, rewind, and bookmark using the Active Timeline tab. Developers have additional functionality with the Timeline History tab that exposes non-active timelines (also known as timeflows).

8.5.1 VDB provisioning wizard

Users can provision VDBs from the DCT UI using the Provision VDB button located under the action button on a VDB details page.

8.5.1.1 Supported database platforms

- **Oracle Multi Instance (Single Tenant)**
- **Oracle Multi Instance (Multi Tenant) with linked vCDB**
 - TDE is supported for Oracle versions greater than or equal to 12.2.
 - Only vCDBs with database versions greater than or equal to 12.1.0.2 can be linked.
 - Auto VDB restart is supported for Oracle versions greater than or equal to 12.1.0.2.
- **Oracle Single Instance (Multi Tenant) with linked vCDB**
 - TDE is supported for Oracle versions greater than or equal to 12.2.
 - Only vCDBs with database versions greater than or equal to 12.1.0.2 can be linked.
 - Auto VDB restart is supported for Oracle versions greater than or equal to 12.1.0.2.
- **MSSQL Multi Instance**
 - Drive letters are not supported and are set to default values.
- **SAP ASE**
- **AppData**
 - AppData is used for connecting with other databases like Postgres, SAP HANA, etc. by uploading the plugin. Refer to the [Provisioning a PostgreSQL VDB](#)⁸² page for more details.
- **vFiles**
 - Refer to the [Unstructured files and app data](#)⁸³ page for more details.

Extending the developer experience capabilities in DCT, users can now provision MSSQL single instance databases or Oracle single instance multi-tenant databases with linked CDB data platforms from the user interface, using an intuitive wizard workflow. Located on the VDB page is a **Provision VDB** button that opens the provisioning wizard. (Note: non admin users will only be able to see provisionable sources (dSources and/or VDBs), environments, and engines to which they are authorized to see and act upon).

8.5.1.2 Step examples

The provisioning wizard will walk through the following steps:

- **Source:** search and select either a dSource or VDB to provision from.

⁸² <https://cd.delphix.com/docs/latest/provisioning-a-postgresql-vdb>

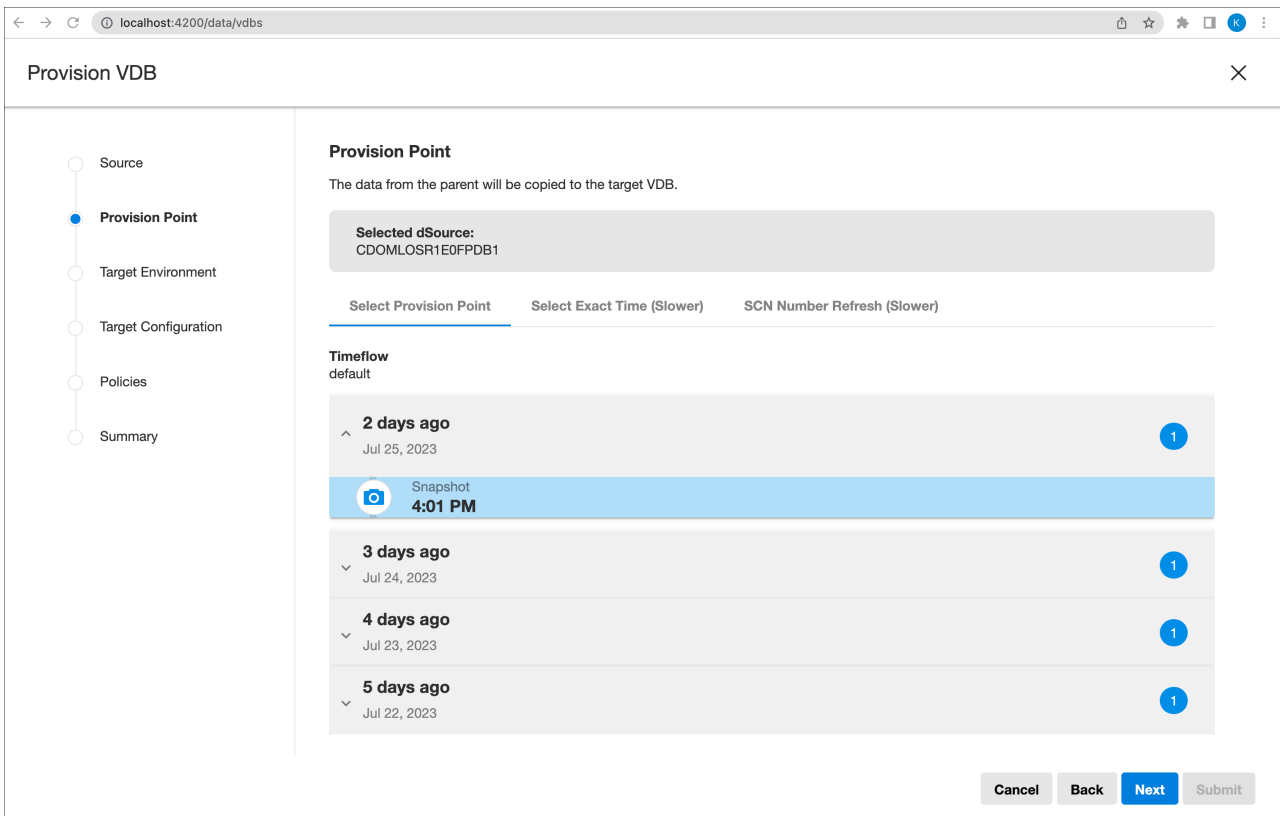
⁸³ <https://cd.delphix.com/docs/latest/unstructured-files-and-app-data>

The screenshot displays the 'Provision VDB' interface. On the left, a vertical navigation menu includes 'Source' (highlighted), 'Provision Point', 'Target Environment', 'Target Configuration', 'Policies', and 'Summary'. The main area is titled 'Source' and instructs the user to 'Select an available Source or VDB from which to provision to your VDB.' Below this, there are two tabs: 'dSources' (active) and 'VDBs'. A search bar is present above a list of sources. The selected source, 'CDOMLOSR1E0FPDB1', is highlighted in blue. To the right of the list, a 'CDOMLOSR1E0FPDB1 Details' panel provides the following information:


- Name:** CDOMLOSR1E0FPDB1
- Version:** Oracle 18.3.0.0.0
- Engine:** KM's Engine
- Group:** Oracle SI MT
- Environment:** Oracle MT 18

At the bottom of the source list, there is a pagination control showing 'Items per page: 25' and '1 - 7 of 7'. At the bottom right of the interface, there are four buttons: 'Cancel', 'Back', 'Next', and 'Submit'.

- **Provision Point:** three options for a provision point, similar to a refresh point.
 - A selected snapshot
 - A specific timestamp (closest snapshot to the timestamp)
 - A location ID/number (closest snapshot to the location number)



- **Target Environment:** shows compatible environments with compatible repositories, and can optionally provide privileged credentials.



- **Target Configuration:** should be prefilled with default configurations. One thing to note here are the tags which are additive when “Include Tags from Parent” is checked, and you wouldn’t immediately see the tags from the parent in the editor. Tags are added when the **Include Tags from Parent** box is checked, you would not immediately see them from the parent in the editor.
 - Users can also select the engine group and register listeners by expanding the advanced section below. In DCT, it is recommended to use tags instead of engine groups.

- The Target Configuration screen will present different options depending on the chosen **Source**.

The screenshot shows the 'Provision VDB' interface with the 'Target Configuration' step selected. The main area contains several input fields for configuration options:

- Mount Base: /mnt/provision
- Oracle Pluggable Database Name: CDOMLOR1E0FPDB1_HE8
- vPDB Name: CDOMLOR1E0FPDB1_HE8
- Oracle Virtual Container Database Name: CDBC DOMLOR1HE8_ECI
- vCDB Name: CDBC_ECI
- Oracle Virtual Container Database SID: CDBC DOMLOECI
- Oracle Virtual Container Database Unique Name: CDBC DOMLOR1HE8_ECI
- Transparent Data Encryption (TDE):

Additional options include 'Include Tags from Parent dSource' (unchecked) and a tag management section with 'Enter Tag Key', 'Enter Tag Value', and a '+ Tag' button. Navigation buttons at the bottom are 'Cancel', 'Back', 'Next', and 'Submit'.

- **vCDB Configure Parameters:** Applicable to Oracle Single Instance (Multi Tenant) with linked vCDB.


The screenshot shows the 'Provision VDB' interface with the 'vCDB Configure Parameters' step selected. The main area displays a table of parameters:

Default Save as New Template

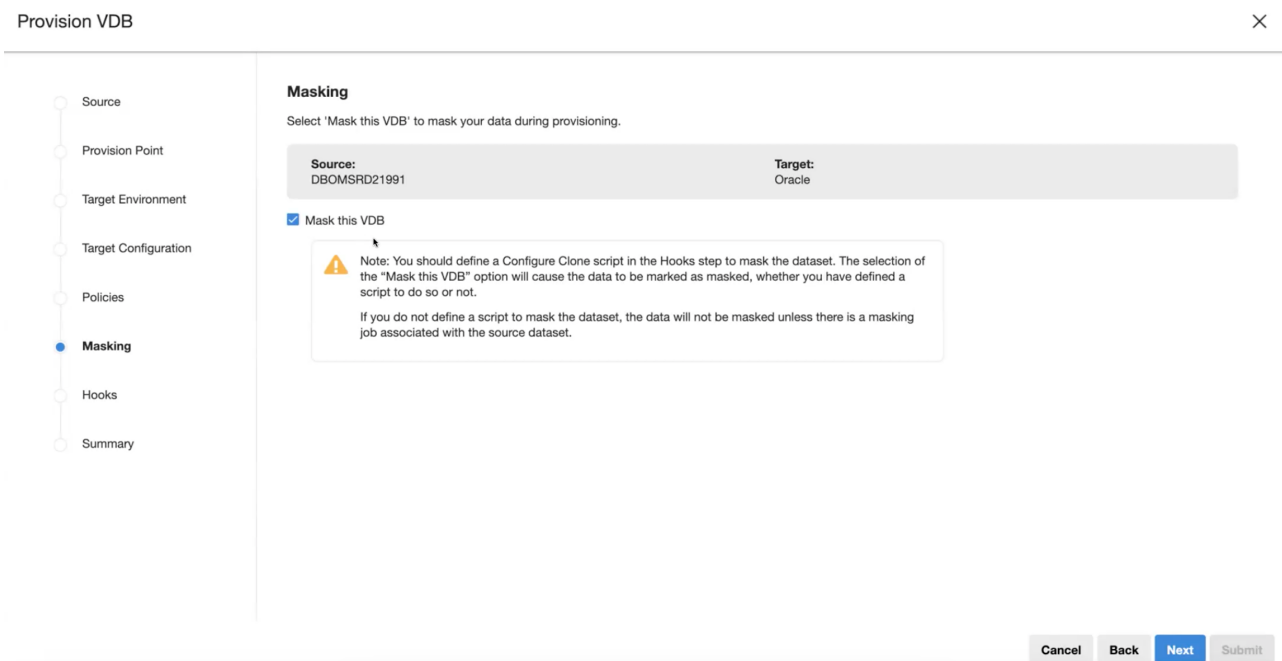
Name	Value
max_pdb	4098
memory_target	1342177280
log_archive_format	%t_%s_%r.dbf
nls_territory	'AMERICA'
compatible	'18.0.0'
processes	300
nls_language	'AMERICAN'

Navigation buttons at the bottom are 'Cancel', 'Back', 'Next', and 'Submit'.

- **Policies:** choose a snapshot policy.



- **Masking:**



- **Hooks:** Delphix users now have the ability to add custom hooks to a VDB at the time of provision (as part of the provision wizard UI) and update them under a VDB detail page, so that they will execute during refresh operations. A Hooks section in the Summary page shows a list of all the hooks configured (if any). Hooks can be modified in the details page after creation.

Hooks

Select what hooks to run during various VDB hook points.

Source: **Alpha Data** Target: **Alpha Environment**

Hook Points

Select hook point to add hooks on the right.

- Configure Clone
- Pre Refresh >
- Post Refresh
- Pre Self Refresh
- Post Self Refresh
- Pre Snapshot
- Post Snapshot
- Pre Start
- Post Start
- Pre Stop

Hooks

Hooks run in order from top to bottom, drag hook cards to re-order.

Custom Created Hook
Remove Edit

Type:
System Shell Command

Description:
Example code to run a script to create AppUser

Hook Selected from Global Hooks
Remove

Type:
System Shell Command

Description:
Example code to run a script to create AppUser

< Back
Next >

- **Summary:** review the selections that have been made. In 10.0 and above, a hooks header has been added to the Summary page that shows a list of all the hooks configured for Hook Operations while provisioning a VDB (if this was configured).

The screenshot shows the 'Provision VDB' Summary page. On the left is a navigation sidebar with steps: Source, Provision Point, Target Environment, Target Configuration, Policies, and Summary (which is selected). The main content area is titled 'Summary' and includes the instruction 'Review the configuration profile for this VDB'. It is divided into two columns: 'Source' and 'Target'.

Source	Target
<p>Source dSource/VDB CDOMLOSR1E0FPDB1_HE8</p>	<p>Environment Oracle MT 18</p> <p>Installation Home /u01/app/oracle/product/18.0.0.0/dbhome_1</p> <p>Mount Base /mnt/provision</p> <p>vPDB Name CDOMLOSR1E0FPDB1_HE8</p> <p>Oracle Pluggable Database Name CDOMLOSR1E0FPDB1_HE8</p> <p>vCDB Name CDBCDOMLOSR1HE8_ECI</p> <p>Oracle Virtual Container Database Name CDBC_ECI</p> <p>Oracle Virtual Container Database SID CDBCDOMLOECI</p> <p>Oracle Virtual Container Database Unique Name CDBCDOMLOSR1HE8_ECI</p> <p>Transparent Data Encryption (TDE) Off</p> <p>VDB Tags No tags added</p>

Cancel Back Next Submit

8.5.1.3 Limitations

TDE and Auto VDB restart are only supported for Oracle version 12.2 or higher.

8.5.2 VDB refresh

8.5.2.1 Overview

The VDB refresh wizard in the Data Control Tower UI offers important engine refresh operations like:


- **Self-refresh** by snapshot, timestamp, or location
 - Refreshes a VDB back to a point in its own history.
- **Refresh to parent** by snapshot, timestamp, or location
 - Data is pulled from the VDB provision parent (the dSource or VDB from which the VDB was provisioned).
- **Refresh to relative** by snapshot, timestamp, or location
 - Allows selection of data from either the origin dSource of the VDB (which could be the immediate parent, or parent of parent, etc.) or any VDB sharing the same origin dSource (which could be parent, child, sibling, or similar relationship).
- **Refresh to bookmark** snapshot
 - Refreshes a VDB to a compatible bookmarked snapshot.




In order to refresh from a dSource or VDB, the account performing the action must have the REFRESH permission on both the VDB being refreshed and the dSource (or VDB) from which the data is being refreshed.

8.5.2.2 User interface

The **self-refresh by snapshot** operation can be initiated as an action from the VDB's own **active timeline** or **timeline history** view, using the action menu for snapshots.



All three refresh types can be performed by opening the refresh wizard on the VDB details page. In the first step, select the refresh source:



Next, select the snapshot or point in time to refresh to:

- For self-refresh, a timeflow from the VDB’s own history must be selected, and then a snapshot within the selected timeflow.
- For parent refresh, a timeflow from the VDB’s provision parent’s history must be selected, and then a snapshot within the selected timeflow.

- For relative refresh, first a relative dSource or VDB must be selected, then a timeflow, then a snapshot.

For all three refresh types, review the summary page once configurations are complete, then click submit:

Summary

Review the following to confirm your action to refresh VDB "Vdbd_3LT" from the selected source.

1 **This action will create a new timeline**
This action will be non-destructive since the existing state will be preserved as an alternate timeline that can be found in Timeline History on the VDB.

<p>Source of Data VDB2SZ</p>	>	<p>Selected VDB Vdbd_3LT</p>
<p>Data Point Snapshot 201</p>		
<p>Origin DSource Time Jul 24, 2019 3:00 PM</p>		



Refreshing happens asynchronously and takes a various amount of time. The DCT UI does not currently show asynchronous job progress or errors, please refer to the engine UI for this function.

8.5.2.2.1 Refresh to Bookmark

Bookmark is available as a **Select Refresh Source** option in the refresh wizard, which allows you to refresh from compatible bookmarks. Choose a bookmark from the list and click **Next**.

Refresh "Vvdb_MIK"
✕

- Select Refresh Source
- Select Bookmark
- Summary

- Select Refresh Source
- Select Bookmark
- Summary

Select Refresh Source

Select the source of the data refresh.

Parent
Set the data state of this VDB to any available point from the parent.

Self
Set the data state of this VDB to any available point on this VDB.

Relative
Set the data state of this VDB to any available point from any relative VDB.

Bookmark
Set the data state of this VDB to any compatible bookmark.

Cancel Back Next Submit

- Select Refresh Source
- Select Bookmark
- Summary

Select Bookmark

Select a bookmark for data refresh.

Search (All) ▾

Bookmark	Creation Date	Tags
my-bookmark-1	Apr 25, 2023 4:33 PM	Add Tags >
my-bookmark-2	Apr 25, 2023 4:34 PM	Add Tags >

Cancel Back Next Submit

You can also refresh from a particular bookmark under the VDB details page, via the Bookmarks tab. Choose a bookmark from the list as the one to refresh from, then click the Actions menu and select **Refresh to Bookmark**.

DATA CONTROL TOWER
Home Data Compliance Insights Admin
admin

VDB

< VDBs **Vvdb_MIK** ⋮

ID: 1-ORACLE_DB_CONTAINER-4

Overview
Active Timeflow
Timeflow History
Bookmarks
Access

Bookmark	Creation Date	Tags	Actions
my-bookmark-1	Apr 25, 2023 4:33 PM	Add Tags >	Actions :
my-bookmark-2	Apr 25, 2023 4:34 PM	Add Tags >	<div style="border: 1px solid #ccc; background-color: white; padding: 5px; width: fit-content;"> <ul style="list-style-type: none"> <li style="padding: 2px 5px;">Refresh to Bookmark <li style="padding: 2px 5px;">Update Bookmark <li style="padding: 2px 5px;">Delete </div>

1 to 2 of 2 ⏪ < Page 1 of 1 > ⏩

8.5.3 Active timelines

8.5.3.1 Active Timeline View

Active Timeline can be found by selecting “details” for any VDB located on the VDB list page in DCT and selecting the “Active Timeline” tab. This view serves as an operations console for any user with the appropriate entitlements to see and act upon the identified VDB granted by the DCT access control system.

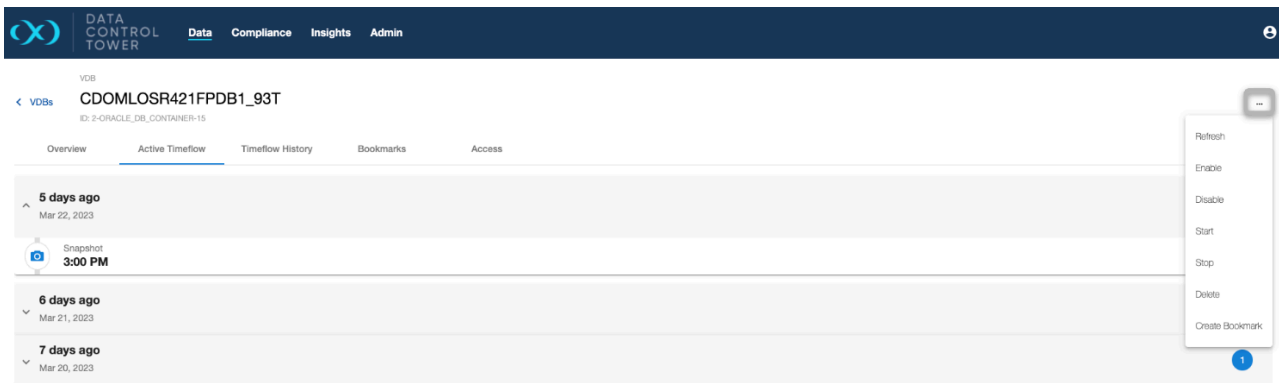
This view shows a vertical timeline-based representation of all actionable points of interest (snapshots) for that VDB. The snapshot list is chronologically grouped over blocks of time to easily identify relevant snapshots to act upon. Once the right snapshot is found, users can access a contextualized action menu by selecting the corresponding ellipsis to the relevant snapshot.



Contextual snapshot menu: By selecting the ellipsis next to the VDB of interest, users can refresh to a snapshot or create a bookmark (a nameable and shareable snapshot reference).

Alternatively, users can perform common Continuous Data actions via the ellipsis in the top right of the Active Timeline screen. This menu will offer VDB-wide, generalized operations including:

- **Refresh:** This is a generalized operation to realign your VDB's timeline with another. More information can be found in the [Refresh a VDB \(see page 212\)](#) article.
- **Start/Stop:** A way to manage target host bandwidth.
- **Enable/Disable:** An administrative operation required for upgrades/migrations.
- **Create Bookmark:** Create a special reference to a snapshot that has a custom name and retention period.



8.5.3.2 Additional notes

- **Refresh using the DCT UI**

Data Control Tower has consolidated all contextualized time-based operations (e.g. refresh, rewind, rollback, restore, reset, etc.) across Continuous Data and Developer products into a general "Refresh". By clicking refresh, users will be prompted on what timeline they would like to align: Parent, Self, or Relative. From there, users will be taken to a wizard that will give relevant timeline options and points in time to perform the refresh.

- **Start/Stop**

Starting and Stopping are geared toward bandwidth management in a target environment. Stopping a VDB will place it in stasis such that it can't be accessed, but also won't consume bandwidth. Starting it back up will re-enable it for regular activity.

- **Enable/Disable**

Enabling and Disabling are geared toward administrative operations 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. Re-enabling the VDB will restore those configuration files.

- **Create Bookmark**

Bookmarks serve as a human-referenceable representation of time that can work for a single VDBs or VDB groups. Bookmarks also have the capability to be shared (**Refresh to relative** in the [Refreshing a VDB \(see page 212\)](#) article). From the top level menu, users will be able to name the reference and set a unique retention period.

Bookmarks do not appear in the Active Timeline, to access existing bookmarks, users will have to navigate to the "Bookmarks" tab on the VDB detail view.

8.5.4 Timeline history

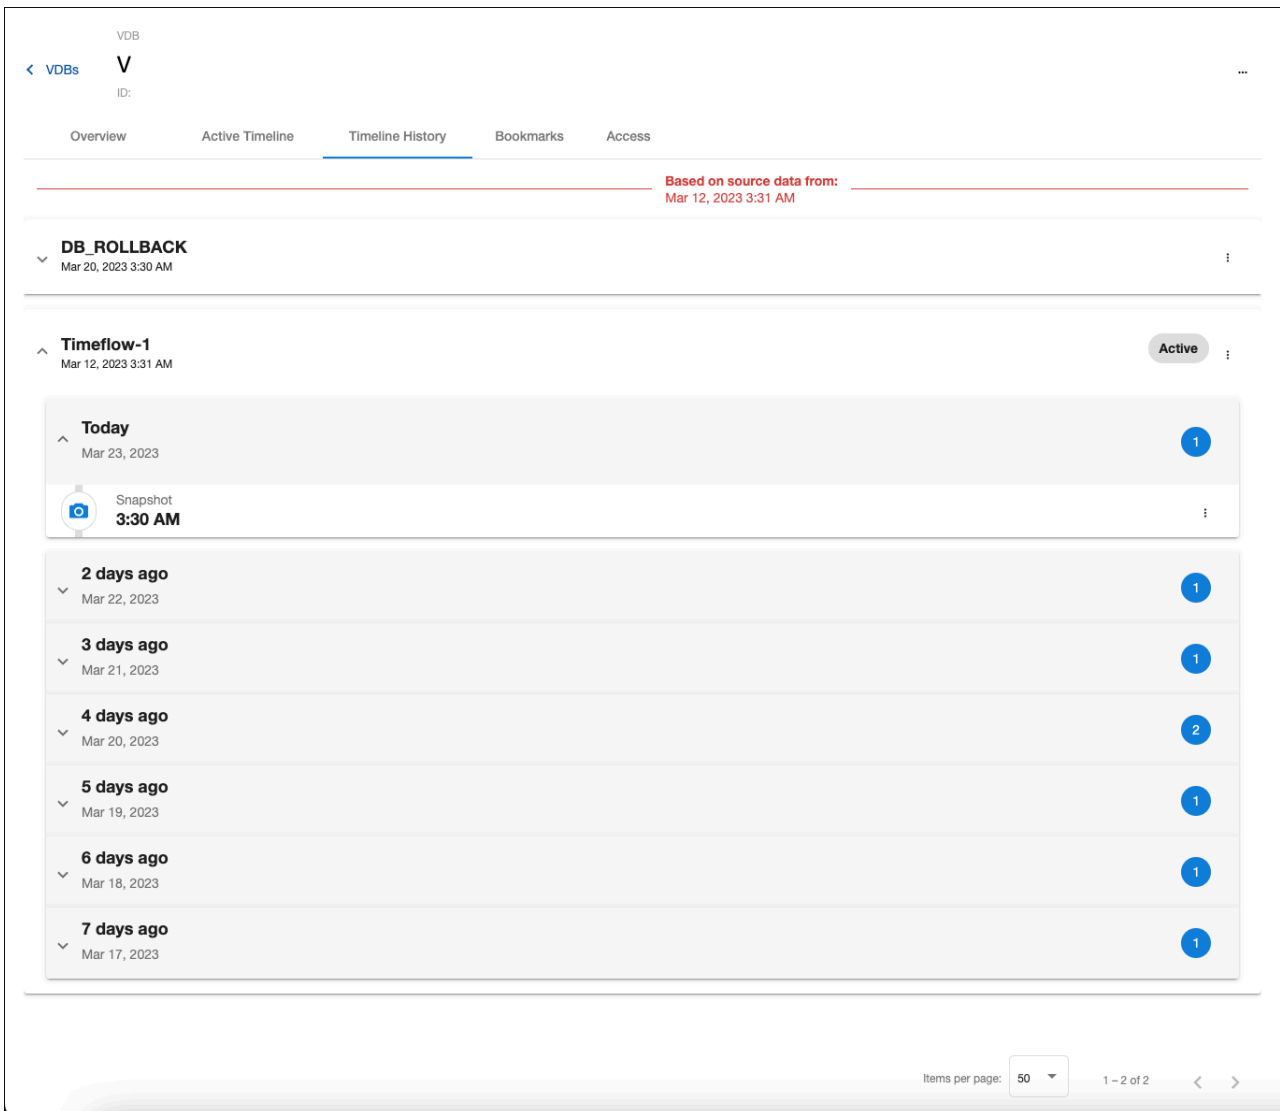
8.5.4.1 Timeline History view

The timeline history view can be found by selecting "details" for any VDB located on the VDB list page in DCT and selecting the "Timeline History" tab. This view serves as a developer-centric console that shows the complete history of a VDB including non-active timelines, which is a critical resource for developers as these can contain information like past test results. Using this page, developers can curate and access the

complete chronology of their testing efforts with operations such as renaming individual timeflows, make active, refresh to snapshot, and create bookmark to drive organization and access.


8.5.4.2 The Timeline History user interface

If entitled via the DCT Access Control system, developers can see and act upon VDBs in the VDB list view using the timeline history tab under "details". The timeline history UI is comprised of a vertical timeline-based representation of all actionable points of interest (snapshots) for that VDB. The snapshot list is chronologically grouped over blocks of time to easily identify relevant snapshots to act upon. Once the right snapshot is found, users can access a contextualized action menu by selecting the corresponding ellipsis to the relevant snapshot.



8.5.4.3 Non-active timelines

Non-active timelines (often referred to as “timeflows” in Self-Service) are a critical aspect of how Delphix Virtualization works and an important Developer Tool. Whenever a time-based operation takes place, the previous timeline (and associated data) becomes non-active and a fresh timeline takes its place.



i Non-active timeline creation from "Self-Refresh"

In this scenario, a developer has performed a “self-refresh” (formerly referred to as a “rewind” or “rollback”). In refreshing Snapshot B, the developer has created a new active timeline that represents a clean slate starting with the data state of Snapshot B. If the developer wants to refer back to the past results that are represented on the non-active Timeline, he or she can activate the past timeline by hitting the “Make Active” button in the DCT User Interface or by refreshing to a point in time by referencing the non-active Timeline’s ID.

For developers, having the ability to catalog (name and tag) and reference past timelines is a critical aspect of application development such as performing ad hoc code validation or manual testing. The timeline history UI in DCT provides a home for single VDB visibility of all accessible timelines (note: timeline availability is controlled through capacity management and snapshot retention policies).

8.5.4.4 Time concepts within the Timeline History Tab


DCT’s new user interface highlights different notions of time within Delphix (e.g. dSource, VDB, and VDB lineages) and how they relate to one another. These time-based relationships are exposed in the Timeline History view to drive accurate testing for developers.

The “Based on dSource time” designation helps to drive awareness of the relationship between a VDB and its dSource provision point (in most cases this will equate to the production database’s state at that time), which is helpful for use cases such as capturing data from a meaningful event. As a VDB refreshes to newer

snapshots on a dSource, those changes of data state are grouped separately on the Timeline History view as they represent completely different data.

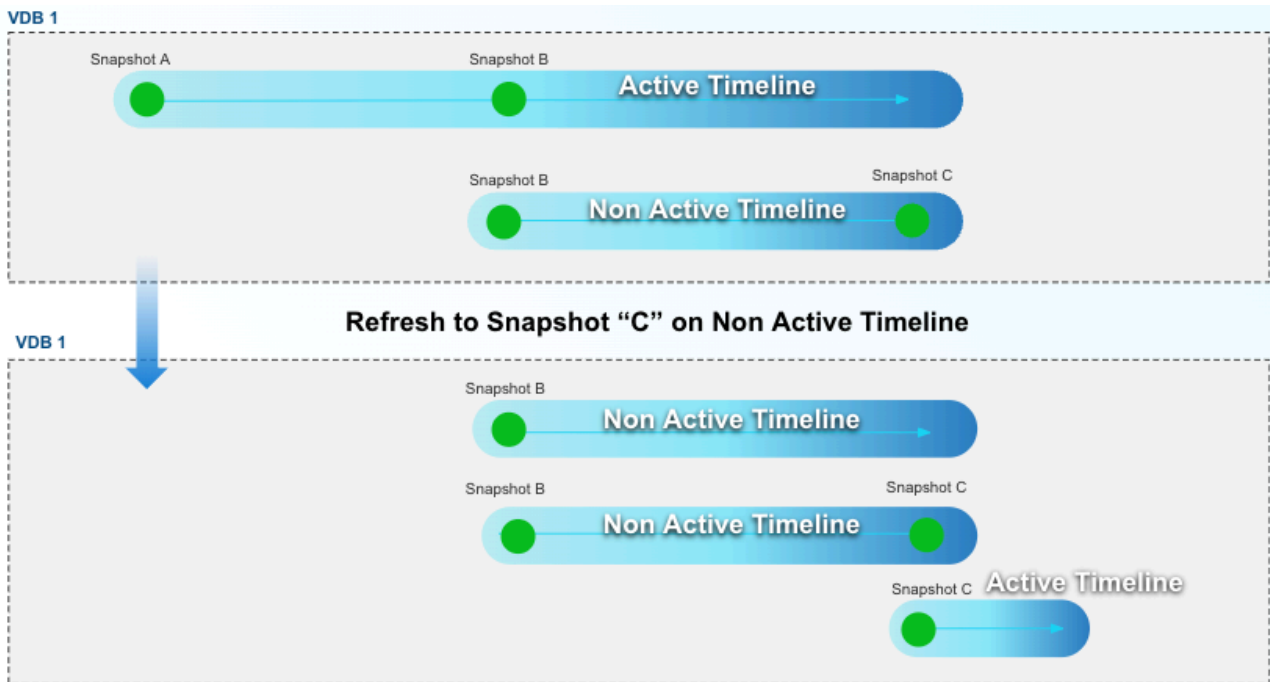
8.5.4.5 Timeflow operations

“Make Active”, a developer can reference any past timeflows and data represented therein by making a timeflow “Active”. This can be done by simply clicking on the timeflow of interest and selecting “Make Active”.



8.5.4.5.1 Refresh to snapshot

Included within each timeflow are snapshots that can be accessed via a dropdown menu under each timeflow. If given the right permission, a developer can refresh the VDB to that snapshot. While this may seem similar to “Make Active”, there are major underlying differences as a “Refresh” will reprovision the VDB based on that point in time, whereas, “Make Active” simply changes the reference to different blocks of the underlying storage.




8.5.4.5.2 Create bookmark

On a snapshot will enable developers to convert that snapshot into a [DCT Bookmark](#) (see page 178) that enables developers to assign a name and special retention policy to that time reference. This also enables stronger collaboration between developers as bookmark references can be used for a sibling refresh operation.

8.5.4.5.3 Rename timeflow

Developers can curate time flows to correspond to tests such that they can easily reference results data. This action can be performed by clicking the ellipsis on the time flow to reach a drop-down menu. This menu includes the option to "rename" which will open a dialog box to input a new name. The input is prefilled with the current Timeflow name by default. Saving the dialog kicks off a job to rename the timeflow. The changes may take a short time to be reflected in the UI.



8.5.4.5.4 Delete timeflow

Developers have the option to delete time flows via the UI by selecting the ellipsis next to the corresponding time flow and selecting "delete". When clicked, this opens a confirmation dialog asking if the user wishes to delete the timeflow. If the dialog is confirmed, a job is kicked off to delete the timeflow. Note: This menu item is disabled for the currently-active timeflow.



8.5.4.6 API documentation

While the DCT UI provides a visual aspect to timeflow chronology and grouping, the DCT API has all of the necessary instrumentation to integrate these concepts into automation. The formal documentation can be found via DCT's swagger docs or the developer resources documentation section. This portion of the timeflow documentation will provide some general examples:

Get a list of a VDB's timeflows

This simple call will enable the listing of the entire timeflow roster underneath a VDB enabling a developer to take inventory and action.

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/timeflows/search?limit=50&sort=id' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "filter_expression": "dataset_id eq '\''1-ORACLE_DB_CONTAINER-2'\''"'
  }'
```

The UPDATE timeflow API can be used to change a timeflow's name

Once the appropriate timeflow is found, developers can rename the timeflow to align with a naming scheme or other relevant designations such as a compatible code branch.

```
curl -X 'PATCH' \
  'https://<APPLIANCE_ADDRESS>/v3/timeflows/1-ORACLE_TIMEFLOW-7' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
```

```
-d '{
  "name": "Latest on refresh state"
}'
```

List all the snapshots for a timeflow

Developers can then use the snapshot API to find all relevant "actionable" points in time on the identified timeflow. The below example shows how to only expose snapshots on a particular timeflow. From there, developers can feed the relevant snapshot ID into refresh or provision endpoints or create a bookmark reference from that snapshot.


```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/snapshots/search?limit=50&sort=id' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
  "filter_expression": "timeflow_id eq '\''1-ORACLE_TIMEFLOW-7'\'''"
}'
```

8.5.5 Locking and unlocking a VDB

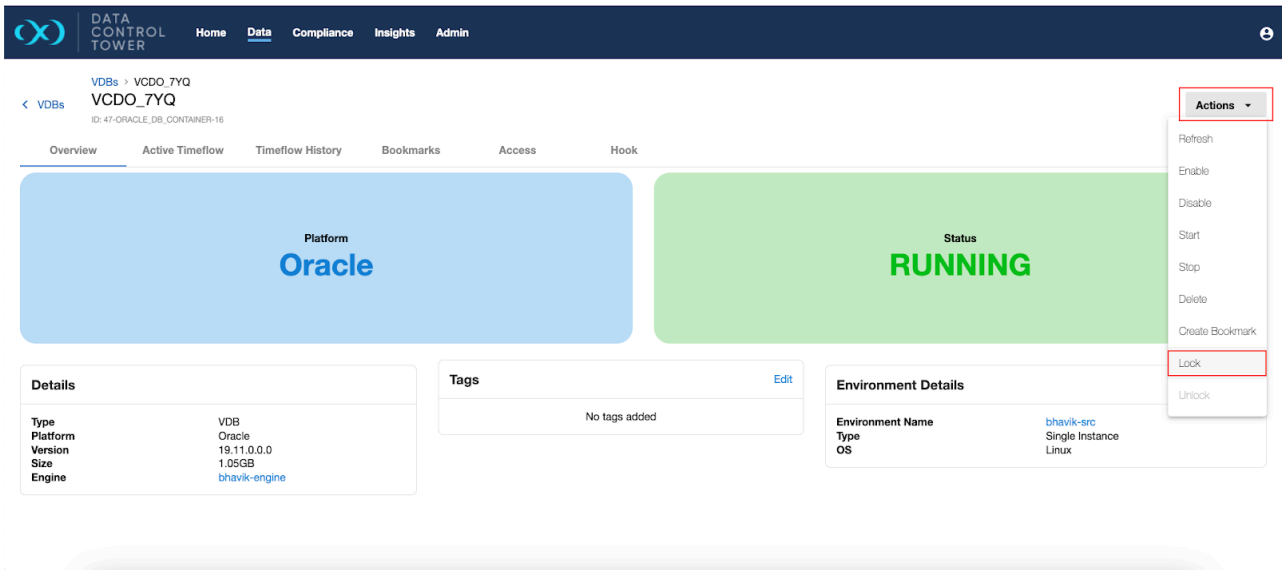
This feature allows users with **LOCK** and **UNLOCK** permission to lock and unlock a VDB, similar to the lock/unlock feature available via Self-Service containers on the Continuous Data Engine.

The VDB can be locked by a user with **LOCK** permissions for that particular VDB. Once a VDB is locked, all of the VDB operations (Refresh, Start, Stop, Delete, and Create Bookmark) are disabled for all other users.


The user needs the **UNLOCK** permission on the VDB to perform an unlock operation. Thus, only the user that locked the VDB **OR** a user with the **FORCE_UNLOCK** permission can unlock it. The **FORCE_UNLOCK** permission is only granted to the **admin** system role by default.

-  Users cannot lock a VDB with refresh policies assigned to it or have any active engine job running at the time of locking.

The VDB details page has the options to lock and unlock a VDB under the Actions menu. All VDBs are in the unlocked state by default, thus, the only option available at first is to lock the VDB.



A warning message is shown on the VDB detail page if the VDB is locked by a user.



Select the Actions menu on a locked VDB to show the unlock action.

The screenshot shows the Data Control Tower interface for a VDB named VCDO_7YQ. The top navigation bar includes Home, Data, Compliance, Insights, and Admin. Below the navigation, there are tabs for Overview, Active Timeflow, Timeflow History, Bookmarks, Access, and Hook. A notification banner states: "VDB is currently locked. Account: 1 has locked this VDB preventing any refresh actions." The main area features two large cards: "Platform Oracle" (blue) and "Status RUNNING" (green). On the right, an "Actions" dropdown menu is open, listing options like Refresh, Enable, Disable, Start, Stop, Delete, Create Bookmark, Lock, and Unlock. The "Unlock" option is highlighted with a red box. Below the main cards, there are sections for "Details", "Tags", and "Environment Details".

8.5.5.1 Locking and unlocking a VDB via API

DCT also provides a capability for an account to lock the VDB on behalf of another account, but this is an API only feature. For an account to be able to lock the VDB on behalf of another account, it needs to have the `LOCK_FOR_OTHER_ACCOUNT` permission on the particular VDB.

Formal documentation around the API signature and related payloads can be found via [DCT API documentation](#)⁸⁴ in the **Developer resources** section. Below is an example of an account requesting to lock a VDB on behalf of an account with `account_id` 2.

```
curl --location 'https://<APPLIANCE ADDRESS>/v3/vdbs/<VDB ID or NAME>/lock' \
--header 'Authorization: <APIKEY>' \
--header 'Content-Type: application/json' \
--data '{
  "account_id" : 2
}'
```

8.6 Linking Wizard

8.6.1 Overview

The DCT version 12.0.0 update marks a significant enhancement to data virtualization capabilities in DCT by introducing the Linking Wizard, which guides the user through intuitive steps to create a dSource.

As of this release, DCT 13.0.0 supports the integration of these types of database sources:

- **Oracle** sources

⁸⁴ <https://dct.delphix.com/docs/latest/api-references>

- Staging Push supported
- **MSSQL** sources
 - Staging Push supported
- **ASE** sources
 - PostgreSQL, Db2, MySQL, HANA, and MongoDB
- **AppData** sources

8.6.1.1 Limitations in MSSQL dSource linking

The ability to link a dSource from backups managed by third-party tools like NetBackup or CommVault is not supported for MSSQL databases as of yet. Users can still create dSources from live MSSQL databases.

This is a known limitation and may be addressed in future updates to expand the versatility and applicability of the dSource feature in diverse data management scenarios.

8.6.1.2 Prerequisites

To utilize the dSource creation feature effectively in DCT 12.0.0, certain prerequisites must be met:

- **Continuous Data Engine registration:** A prerequisite for dSource integration is having a Continuous Data Engine already registered to a DCT instance. This setup is crucial for enabling the core functionalities of data virtualization and synchronization.
- **Availability of eligible sources:** It is necessary to have access to eligible database sources that can be linked as a dSource. These sources should not already be functioning as a dSource, to avoid redundancy and ensure optimal utilization of the new feature.

8.6.2 Using the Linking Wizard

This new wizard can be launched from the dSource list page, giving the user the ability to create a dSource from DCT. The steps are outlined below:

8.6.2.1 1. Select Data Source step

Navigate an alphabetically sorted list of data sources, or utilize the search function and pagination to locate and choose the preferred source.

- Upon selecting a source, essential details such as the **Instance**, **Environment Name**, and **Engine Name** are displayed on the right-hand side of the screen. **Database Authentication** will be required.
- To that point, a crucial part of this step includes a database connectivity check. The **Validate** button verifies your connection, ensuring a secure and authenticated progression to the next stages.

The screenshot shows the 'Create and link dSource' wizard with the following details:

- Navigation:** Select Data Source (active), dSource Configuration, Data Management, Policies, Hooks, Summary.
- Select Data Source:** Select an available Data Source to connect and create a dSource.

Delphix_Admin MSSql 11.3.6523.0
Delphix_Admin MSSql 13.1.4001.0
Ferrari MSSql 13.2.5233.0
Ferrari MSSql 13.2.5233.0
Impala MSSql 11.3.6523.0
Impala MSSql 11.3.6523.0
Impala MSSql 11.3.6523.0
- Configuration for Selected Source (Ferrari):**
 - Instance:** SQL2016
 - Name:** Ferrari
 - Database Version:** 13.2.5233.0
 - Environment Name:** (empty)
 - Data Type:** MSSql
 - Engine:** (empty)
 - Database Authentication:**
 - Database User
 - Environment User (Environment User)
 - Domain User with Password Credential
 - Environment User:** qa-ad\delphix
- Buttons:** Validate, Cancel, Back, Next, Submit.

8.6.2.2 2. dSource Type step

In this step, select the nature of the dSource being linked. The options available include:

- **Source Based dSource:** linked directly to a source database.
- **Oracle Staging Push:** linked to a user-managed Oracle staging database.
- **MSSQL Staging Push:** linked to a user-managed MSSQL staging database.

Following the selection of the dSource type, users must then choose the associated engine from the dropdown menu, which lists the available engines that have been registered in the DCT environment. This setup dictates the subsequent configuration steps to ensure the linked dSource is correctly aligned with the user's infrastructure and database management strategy.

Create and link dSource
✕

- dSource Type
- dSource Configuration
- Data Management
- Policies
- Hooks
- Summary

dSource Type

Select a dSource type.

Source Based dSource
dSource linked to a source database

MSSQL Staging Push
MSSQL dSource linked to customer managed staging database

Oracle Staging Push
Oracle dSource linked to customer managed staging database

Select Associated Engine
 Engine

Cancel Back Next Submit

8.6.2.3 3. dSource Configuration step

Name your dSource here. By default, the dSource name mirrors the name of the source, but it can be modified. Choose the **Target Group** within the engine to which the new dSource will be added. Add **Notes**, and a **Tag Key** and **Tag Value** pair, if needed. You must also configure any necessary settings required, based on the source.

Create and link dSource
✕

- Select Data Source
- dSource Configuration
- Data Management
- Policies
- Hooks
- Summary


dSource Configuration

+ Tag

Cancel Back Next Submit

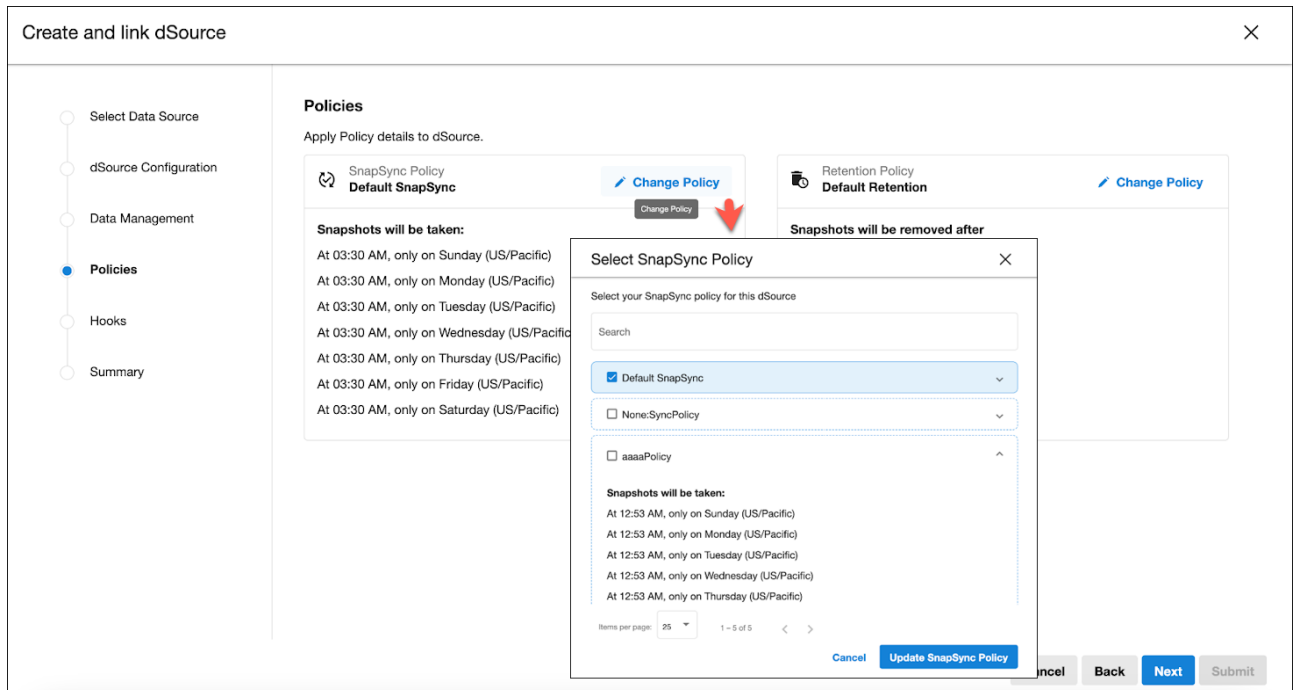
8.6.2.4 4. Data Management step

The step contains various dSource configurations for optimal data management. The wizard fetches and pre-populates the default values for various data fields, saving you time and reducing manual entry errors.



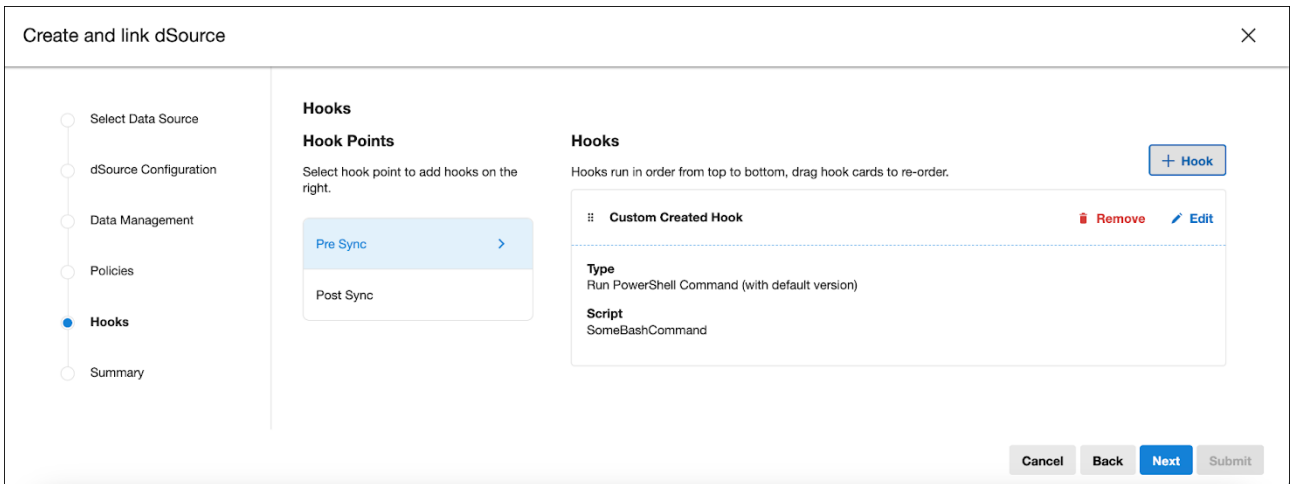
8.6.2.5 5. Policies step

Tailor the SnapSync and Retention policies according to your requirements. Default policies are selected initially, but these settings can be modified through the **Select Policy** dialog.



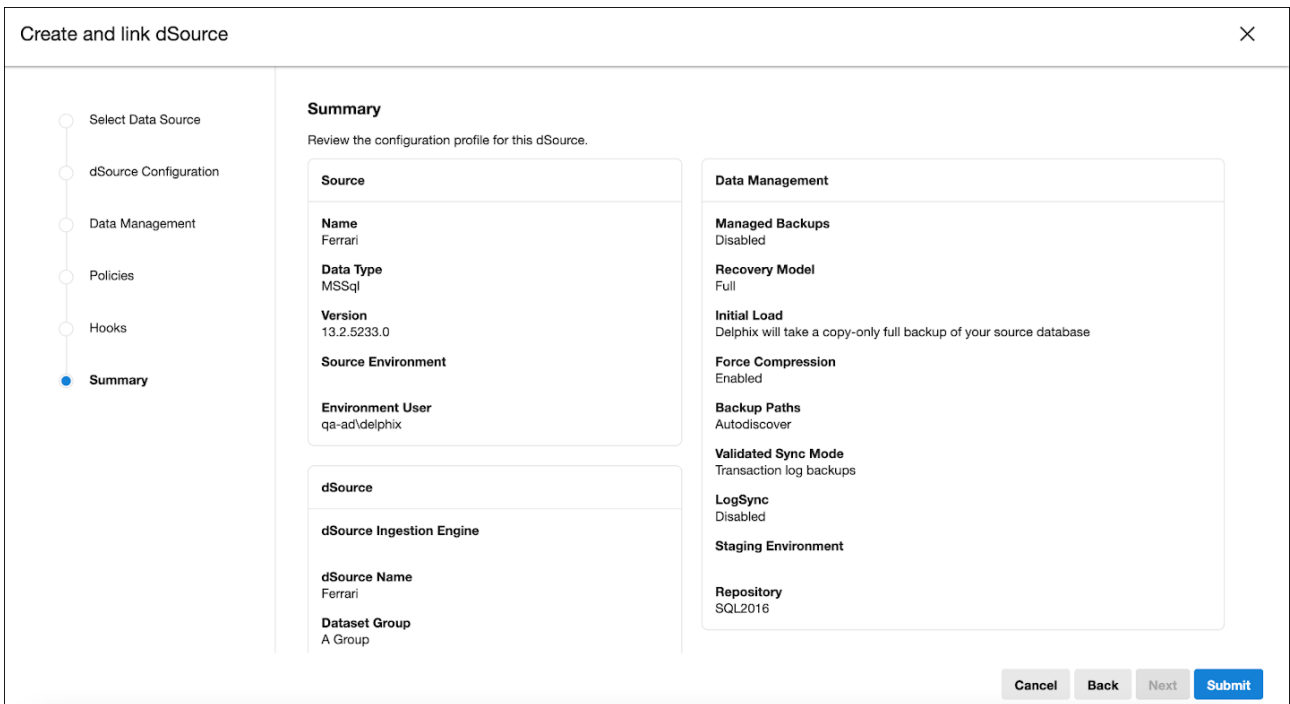
8.6.2.6 6. Hooks step

Enhance the dSource with custom created Hooks. Use the **+ Hook** button to define the scripts that will execute at various stages of the process, offering greater control and customization over the dSource creation workflow.



8.6.2.7 7. Summary step

Before finalizing, a comprehensive summary of all your configurations is presented. This overview allows you to review and confirm that every detail aligns with the intended setup, ensuring accuracy and completeness. **Submit** the configuration to finish linking your database.



8.6.3 Adding an AppData dSource

The steps to add a dSource in DCT, as outlined in the [Linking wizard](#)⁸⁵ page, should be still be followed. However, this page covers details specific to adding an AppData dSource to DCT. AppData dSources allow for efficient handling and integration of application data.

Before initiating the linking process, ensure that an AppData type source is already added in DCT.

8.6.3.1 Steps effected during the Linking process

8.6.3.1.1 Data Source step

1. **Select AppData source:** Initiate the dSource Linking wizard and select the desired AppData source. This selection is crucial in determining the linking procedure.

The screenshot shows the 'Create and link dSource' wizard. The sidebar on the left indicates the current step is 'Select Data Source'. The main area is titled 'Select Data Source' and contains a search bar with 'appdata' entered. Below the search bar is a list of sources, with the first one selected: 'Source 1 nix_staged_python'. The 'Source 1 Details' panel on the right shows the following information:

- Name: Source 1
- Environment Name: (empty)
- Data Type: AppData
- Engine: (empty)
- Environment User: oracle
- Add delay: Add delay
- Add 30 second delay to operation: None
- Add Shell Pause: None
- Add 60 second pause in selected method code: None
- Add 60 second pause in selected method inside a run powershell call: None
- dArray: Direct Source Array

At the bottom right of the main area is a '+ Add' button. At the bottom of the wizard are 'Cancel', 'Back', 'Next', and 'Submit' buttons.

2. **Dynamic UI form:** For an AppData database, a toolkit associated with it contains a `LinkedSourceDefinition`. This definition, essentially a JSON schema, is rendered as a Dynamic UI form alongside the Source list. It outlines the necessary steps for linking sources of this type.
 - **Direct AppData source:** If selecting a direct AppData source, there's an option for adding advanced settings, provided the `LinkedSourceDefinition` has `usesGrandfatheredAppDataProperties` set to true.

⁸⁵ <https://dct.delphix.com/docs/latest/linking-wizard>

Create and link dSource
✕

- dSource Type
- Select Data Source**
- Advanced settings
- dSource Configuration
- Policies
- Hooks
- Summary

Select Data Source

Select an available Data Source to connect and create a dSource

source 1
nix_staged_python
Source 1
nix_staged_python
Source 2
nix_staged_python
Source 2
nix_staged_python
Source 2
nix_staged_python
Source 2
nix_staged_python
myDB
Unstructured Files
source sourabh
Second_Plugin
source sourabh
Second_Plugin
source sourabh
Second_Plugin

Items per page: 1 - 13 of 13

myDB Details

Name
myDB

Environment Name

Data Type
AppData

Engine

Environment User

Advanced settings

Create and link dSource

Advanced settings

Paths to Exclude

Path1

Path

Path2

Paths of Symlinks to Follow

Link1

8.6.3.1.2 Data Management step

- This step is conditional and appears if:
 - a. The source is a staged type.
 - b. The source, whether direct or staged, requires snapshot parameter configuration.
- **Snapshot Parameters:** These are determined by the toolkit's `snapshotParametersDefinition` and are presented as a Dynamic UI form.
- **Staging Environment selection:** Users can choose the appropriate staging environment and user credentials during this step.

Create and link dSource
✕

- dSource Type
- Select Data Source
- dSource Configuration
- Data Management**
- Policies
- Hooks
- Summary

Data Management

Staging Environment

User

Snapshot Parameters

Parameters for the first snapshot taken on the new dSource.

resync

snapParamArray + Add

Snapshot Parameter Boolean Value
A boolean value.

An enum value.

An integer value.

An Number value.

Cancel
Back
Next
Submit

8.6.3.1.3 Completion

Continue through the linking process and after submission, the process of creating an AppData dSource begins. The new dSource will soon be visible on the dSource page, reflecting the configurations made during the linking process.

8.6.4 Adding an Oracle Staging Push dSource

This page outlines the process for adding an Oracle Staging Push dSource within Delphix Data Control Tower (DCT). This functionality is crucial for managing Oracle databases, particularly in environments where staging operations are pushed to target engines. Integrate Oracle Staging Push dSources into your DCT setup by following the Linking Wizard steps.

8.6.4.1 Prerequisites

Ensure that a target engine, with the source backup already cloned, is registered in the DCT before beginning the linking process.

8.6.4.2 Linking process

8.6.4.2.1 dSource Type step

1. **Select dSource type:** Begin the dSource Linking Wizard and choose the 'Oracle Staging Push' option.
2. **Target engine selection:** From the engine dropdown, select the desired target engine where the source backup is located.

8.6.4.2.2 dSource Configuration step

- **Select dSource configuration:** Depending on the type of Oracle database (CDB, PDB, or single tenant), select the appropriate dSource configuration.
- **Database details:** Based on the previous selection, enter the necessary details that would appear as shown in the screenshot below.

Create and link dSource

×

- dSource Type
- dSource Configuration**
- Data Management
- Hooks
- Summary

dSource Configuration

Database Type

Non Multi Tenant

CDB

PDB

Custom Environment Variables

<small>Key</small> var1	<input type="button" value="Delete"/>
<small>Value</small> var1 value	

- **Staging database parameters:** An option to configure Staging Database Parameters is available. If selected, it adds a new step for entering these parameters.

Create and link dSource ✕

- dSource Type
- dSource Configuration
- Staging Database Parameters
- Data Management
- Hooks
- Summary

Staging Database Parameters

Save as New Template

Table
Text

Name	Value
parameter	parametervalue

Cancel Back Next Submit

8.6.4.2.3 Data Management step

- **Enable LogSync:** This step involves the option to enable log synchronization for real-time data updates and management.

Create and link dSource ✕

- dSource Type
- dSource Configuration
- Staging Database Parameters
- Data Management
- Hooks
- Summary

Data Management

LogSync

Enable (Archive Only)

LogSync is an add-on service that can synchronize the Delphix Engine within seconds of any change to a source database. LogSync requests and stores changes to log files. By coordinating changed log blocks with snapshots from SnapSync, users can provision VDBs from a specific point in time, either an arbitrary time, or a time relative to the last transaction.

Cancel Back Next Submit

8.6.4.2.4 Summary step

- **Review configuration:** Examine the details of the summary, which includes information about the dSource type, target engine, database details, and any additional parameters set in the previous steps, to ensure all configurations are correct and aligned with the requirements of the Oracle Staging Push dSource.

8.6.4.3 Completion

Upon confirming and submitting the details in the Summary step in the Linking Wizard, the process of creating an Oracle Staging Push dSource is initiated. The new dSource will be integrated into the DCT environment and will be available for use once the setup is complete.

8.6.5 Adding a MSSQL Staging Push dSource

This page outlines the process for adding a MSSQL Staging Push dSource within Delphix Data Control Tower (DCT). This functionality is crucial for managing MSSQL databases, particularly in environments where staging operations are pushed to target engines. Integrate MSSQL Staging Push dSources into your DCT setup by following the Linking Wizard steps.

8.6.5.1 Prerequisites

Ensure that a target engine, with the source backup already cloned, is registered in the DCT before beginning the linking process.

8.6.5.2 Linking process

8.6.5.2.1 dSource Type step

1. **Select dSource type:** Begin the dSource Linking Wizard and choose the 'MSSQL Staging Push' option.
2. **Select Associated Engine:** From the engine dropdown, select the desired target engine where the source backup is located.

The screenshot shows a wizard window titled "Create and link dSource" with a close button (X) in the top right corner. On the left side, there is a vertical progress indicator with six steps: "dSource Type" (selected with a blue dot), "dSource Configuration", "Data Management", "Policies", "Hooks", and "Summary". The main content area is titled "dSource Type" and contains the instruction "Select a dSource type." Below this, there are three radio button options:

- Source Based dSource
dSource linked to a source database
- MSSQL Staging Push
MSSQL dSource linked to customer managed staging database
- Oracle Staging Push
Oracle dSource linked to customer managed staging database

 Below the radio buttons is a dropdown menu labeled "Select Associated Engine" with the text "Engine" and a downward arrow. At the bottom right of the wizard, there are four buttons: "Cancel", "Back", "Next" (highlighted in blue), and "Submit".

8.6.5.2.2 dSource Configuration step

- **Configure dSource details:** Enter the dSource Name, Database Name, select a Target Group from the dropdown, and enter Notes or add Tags, if desired.

Create and link dSource

×

- dSource Type
- dSource Configuration**
- Data Management
- Policies
- Hooks
- Summary

dSource Configuration

+
Tag

Cancel
Back
Next
Submit

8.6.5.2.3 Data Management step

- **Staging Environment selection:** Choose the appropriate Staging Environment to host the staging database from the dropdown list.
- **Repository selection:** Select the corresponding Repository version.

Create and link dSource ✕

- dSource Type
- dSource Configuration
- Data Management**
- Policies
- Hooks
- Summary

Data Management

Staging Environment ▾

.co

This environment will host the staging database used for validated sync. Repository

SQL2016 ▾

Cancel
Back
Next
Submit


8.6.5.2.4 Summary step

- **Review configuration:** Examine the summary, which includes information about the dSource type, target engine, database details, and any additional parameters set in the previous steps, to ensure that all configurations are correct and aligned with the requirements of the MSSQL Staging Push dSource.

8.6.5.3 Completion

Upon confirming and submitting the details in the Summary step in the Linking Wizard, the process of creating an Oracle Staging Push dSource is initiated. The new dSource will be integrated into the DCT environment and will be available for use once the setup is complete.

9 Continuous Compliance workflows

 Compliance Engines limit any syncing operations while a profiling or masking job is running. When using DCT job **move**, **execute**, or **migrate** operations, please ensure that the target Compliance Engine is in an idle state. Future Compliance enhancements to DCT will remove this limitation.

With the ability to distribute and run jobs, DCT enables advanced Compliance Engine architectures to be orchestrated and monitored using DCT's real-time, persistent relationships with connected Compliance Engines. When syncing a Compliance Engine, DCT will create references for all Compliance jobs on that Engine. These will show up as unique objects tracked by DCT that can now be leveraged with job move APIs.

9.1 Listing and searching compliance jobs

When a Compliance Engine is registered with DCT, compliance jobs (referred to as MaskingJobs within the DCT API) on the Engine are automatically ingested and presented as DCT MaskingJob objects.

Example of **listing all MaskingJobs**:

```
curl -X 'GET' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>'
```

Example of **searching for OnTheFly MaskingJobs**:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/search' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
  "filter_expression": "is_on_the_fly_masking eq true"
}'
```

With the new job move APIs, DCT can now be used to power two advanced [masking reference architectures](#)⁸⁶: **Software Development Lifecycle (SDLC)** and **Horizontal Scale** architectures. SDLC enables the separation of duties for the development, quality assurance, and production use of masking jobs whereas Horizontal Scale enables the use of a central configuration engine with the movement of jobs to headless compute engines.

⁸⁶ <https://delphixdocs.atlassian.net/wiki/spaces/CC/pages/12192185>

To enable these architectures, DCT has introduced three new operations: Job Copy, Job Execute, and Job Migrate:

- **Copy:** Supports SDLC by copying a job, but maintaining separate references in DCT.
- **Execute:** Supports Horizontal Scale by copying a job, but maintaining the same reference between two copies. DCT will also keep both of these copies in sync.
- **Migrate:** Supports the movement of a single instance from one engine to another.

9.2 Consolidated operations (intelligent syncing)

DCT has simplified the set of operations required to move a job and its dependencies. Previously, orchestrating movement of jobs required three separate API calls: Job Sync, Global Object Sync, and Credentials Update (on the newly created job). DCT has now consolidated all three of these operations into each of the job move APIs. In addition, if two jobs are held in sync (see [Job Execute \(see page 250\)](#)), DCT will auto update synced jobs whenever one of those jobs has been modified (i.e. updated rule set, new algorithms, etc.).



In order to transfer connector credentials with a job as part of the job move, you will need to associate those credentials using the connector credentials API. See sample code below on how to update credentials.

Example of updating a MaskingJob with connector credentials:

```
curl -X 'PATCH' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead' \
  \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "connector_username": "user123",
    "connector_password": "password123"
  }'
```

9.3 Managing engines (Continuous Compliance)

DCT provides a near real-time list of all connected Continuous Compliance engines and lists them in an aggregate view. From the below screen, Delphix administrators can easily view and manage their engine connections.

Compliance Engines Connect Engine

Overview of all compliance engines in your Delphix ecosystem.

Search Search (All) ▾

Status	Name/UUID ↑	Type	Running Jobs	Masking Usage/Total	Masking Available Cores	Tags
● Online	Test Engine 2	Masking 9.0.0.0	0	0% 2.00GB	4	View tags (1) > View >
● Online	Test Masking	Masking 9.0.0.0	0	0% 2.00GB	4	Add Tags > View >

From this screen, administrators can manage engine connects via the “Connect Engine” button on the top right corner. By clicking this button, the below window will appear asking for connection details.

ⓘ DCT will access the engine as a registered user and, as detailed in the Deployment section, requires both a username and password as well as admin-level access to the engine. For compliance engines, select "Masking" type when registering an engine.

Connect Engine ×

- Engine Details
- Authentication
- Security
- Tags
- Summary

Engine Details

Name:

Hostname:


Choose Engine Type

Virtualization

Masking


9.3.1 Engine overview

Individual engine details can be seen and acted upon by clicking down on a particular engine detailed view. Once clicked, users will be sent to an "overview" tab that provides relevant metadata related to the engine.



9.3.2 Engine-based operations access

Users are able to audit which users have access to this particular engine, what access group they belong to, and the associated permissions that each user has on this engine. Admins are able to click on the "View" button to access further details under the access control screen related to that specific user.



9.4 Compliance jobs

- [Job UI \(see page 244\)](#)
- [Copy job \(see page 247\)](#)
- [Execute job \(see page 250\)](#)
- [Migrate job \(see page 252\)](#)
- [Delete job \(see page 253\)](#)

9.4.1 Job UI

9.4.1.1 Global compliance jobs list

When connected to a Continuous Compliance engine, DCT will sync and create references to every compliance job on the engine. All of those job references can be found in the global compliance list, which is a taggable, filterable, sortable, and searchable list of all compliance jobs across a connected Delphix ecosystem.

The screenshot shows the 'Compliance Jobs' page in the Data Control Tower interface. The page has a dark blue header with the 'DATA CONTROL TOWER' logo and navigation tabs for 'Data', 'Compliance', 'Insights', and 'Admin'. A left sidebar contains 'Compliance Engines' and 'Compliance Jobs'. The main content area is titled 'Compliance Jobs' and includes a descriptive sentence: 'Compliance jobs are the configurations that enable you to mask data on your data sources.' Below this is a search bar with a 'Search (All)' dropdown. A table lists six compliance jobs, each with a 'View >' link.

Name/ID ↑	
Drupal 9 Custom App Masking 020e35d2-7d60-482c-908b-7825e89ed3e7	View >
Drupal 9 Custom App Masking 08edd39c-bbe2-4296-9e10-ac255fdc8722	View >
M1-sCRM-Dev 473b6938-106a-4969-bc53-d756f33845fa	View >
M1-sCRM-Dev 4ed7691e-cab8-48aa-a14a-62339d9a0534	View >
Mask anagrafe 36bfd517-30d0-400f-b950-5ea8985426eb	View >
Mask anagrafe bbfd11e9-3cc9-4998-bc40-a624f0a15241	View >

9.4.1.2 Compliance job overview

Individual compliance job details can be seen and acted upon by clicking down on a particular compliance job detailed view. Once clicked, users will be sent to an "overview" tab that provides relevant metadata related to the VDB.

Compliance Job

Compliance Jobs **Drupal 9 Custom App Masking** ID: 020e35d2-7d60-482c-908b-7825e89ed3e7

Overview Access History

Job Status : SUCCEEDED Run Time : 22secs
Started at 3:09 AM

- ✓ Initializing
- ✓ Collecting Configurations
- ✓ Preparing
- ▶ Running Pre-execution Custom Driver Tasks
- ✓ Starting
- ▶ Running PreSQL Scripts
- ▶ Running PostgreSQL Scripts
- ▶ Running Post-execution Custom Driver Tasks
- ✓ Collecting Information
- ✓ Job Completed

Details

Created
May 16, 2022 3:54 AM

Rule Set
Drupal 9 Custom app

Tags [Edit](#)

9.4.1.3 Access tab

Users are able to audit which users have access to this particular compliance job, what access group they belong to, and the associated permissions that each user has on this job. Admins are able to click on the "View" button to access further details under the access control screen related to that specific user.

Compliance Job

Compliance Jobs **Drupal 9 Custom App Masking** ID: 020e35d2-7d60-482c-908b-7825e89ed3e7

Overview Access History

User	Access Groups	Permissions	
	1 >	CANCEL COPY DELETE	View >
John Smith johnsmith@company.com	John Smith >	CANCEL COPY DELETE	View >

Masking Jobs

- CANCEL Copy Delete
- Execute Manage Tags MIGRATE
- Read Update

9.4.1.4 Execution history tab

The compliance jobs page also includes job execution history. Execution information will be recorded and displayed on this tab, and includes the status (success, running, failed), run time, submit, and end timestamps, and engine on which it ran. Additionally, clicking the "details" button for an execution will display its report.

Compliance Job

< Compliance Jobs **ComplianceJob1** ...

726617c0-6fe8-49e1-8b25-4ae46339864c

Overview Access History

Search Search (All) v

Run Time	Submit Time	End Time ↓	Engine Name	
✘	May 3, 2023 6:44 PM	May 3, 2023 6:44 PM	em-masking	Details >
✔	May 3, 2023 6:44 PM	May 3, 2023 6:44 PM	em-masking-6013.dc...	Details >
✔	May 3, 2023 6:44 PM	May 3, 2023 6:44 PM	em-masking-6013.dc...	Details >

1 to 3 of 3 < Page 1 of 1 >

9.4.1.5 Execution details

The execution details view includes a list of execution events and the execution log. This is particularly useful when troubleshooting failed executions.

- Details for successful executions may also have events and logs which include relevant information, such as warnings.

Error Report

i

Execution event Information

This section shows exceptions that occurred during the job run. More detailed information may be available in Masking logs.

Execution Events

Event	Cause	Description
JOB_ABORTED	UNHANDLED_EXCEPTION	Listener refused the connection with the foll...

Logs


```

2023-05-05 04:44:13, [thread] INFO com.dmsuite.dmsApplicator.masking.XMLGenerator executeMarshalling - Generate request xml started successfully.
2023-05-05 04:44:13, [thread] INFO com.dmsuite.dmsApplicator.masking.XMLGenerator executeMarshalling - Generate Request xml done successfully.
2023-05-05 04:44:13, [thread] INFO com.dmsuite.dmsApplicator.masking.transformation.MaskingMarshalling createKettleXML - Generate Transformation XML started successfully
2023-05-05 04:44:13, [thread] SEVERE com.dmsuite.dmsApplicator.masking.transformation.MaskingMarshalling handleJobGenerationException - An exception occurred during job generation
java.sql.SQLException: Listener refused the connection with the following error:
ORA-12505, TNS:listener does not currently know of SID given in connect descriptor


```

OK

9.4.2 Copy job

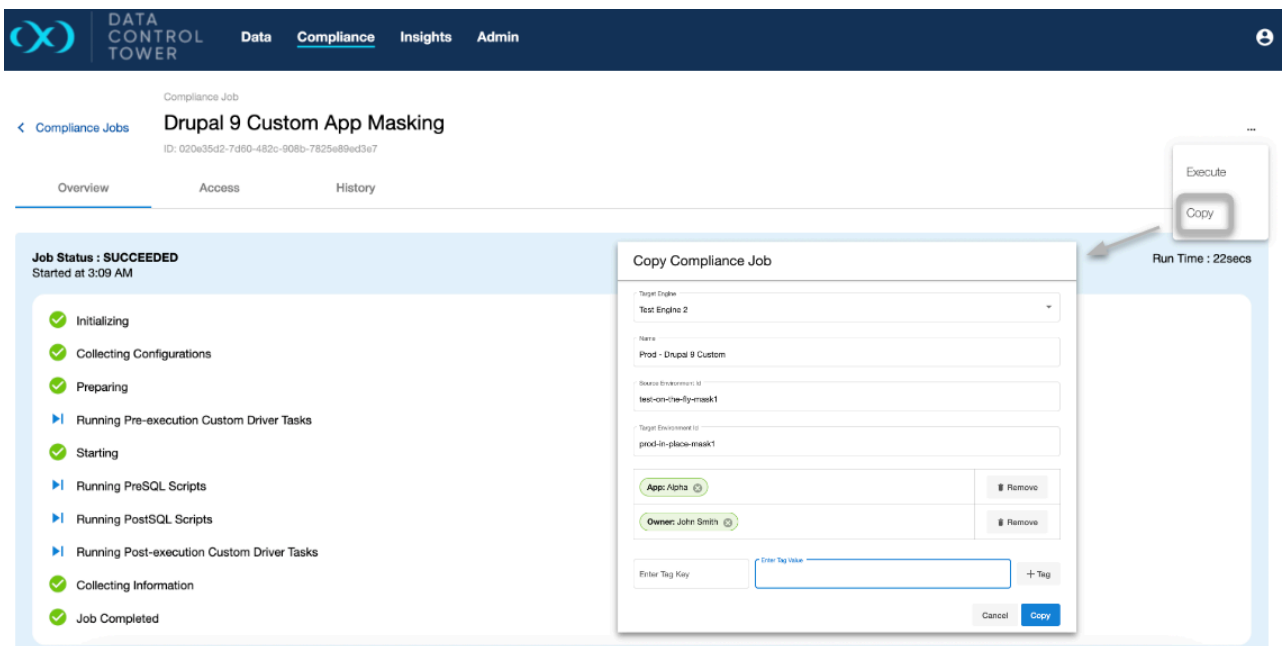
 Compliance Engines limit any syncing operations while a profiling or masking job is running. When using DCT job **move**, **execute**, or **migrate** operations, please ensure that the target Compliance Engine is in an idle state. Future Compliance enhancements to DCT will remove this limitation.

The Masking Job **Copy operation** creates a duplicate of a job with a separate reference for that new copy. This operation supports SDLC workflows as DCT will maintain unique references for each instance of a masking job, enabling them to be managed independently.



9.4.2.1 User interface documentation

Job copy can be run via the DCT UI by accessing a compliance job's detailed view and selecting the ellipsis in the top right corner and clicking on "copy". This will open a window to select the target engine, the new name of the transferred job, source and target environment details, and relevant tags.



9.4.2.2 API documentation

For input, the user must specify the target engine along with the environment on the target engine that the job will be copied onto. The engine and environment pair is what uniquely identifies a copy of the Masking

Job. Calling the COPY API against the same target engine and environment effectively serves as a re-sync and does not create a new DCT MaskingJob entity.

Example of **copying a MaskingJob to engine with ID 2 and environment named 'prod-env'**:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead/copy' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "target_engine_id": "2",
    "target_environment_id": "prod-env"
  }'
```

MaskingJob sync will not copy connector credentials to another engine. In order to make a copied job executable outside of DCT, the credentials must be set on the Connector itself. The connectors for a MaskingJob can be searched for, updated, and tested directly via DCT.

Example of **listing connectors for a MaskingJob**:

```
curl -X 'GET' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead/connectors' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>'
```


Example of **updating a connector's credentials**:

```
curl -X 'PATCH' \
  '<https://<APPLIANCE_ADDRESS>/v3/connectors/2-DATABASE-23>' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "username": "USER123",
    "password": "password123"
  }'
```

Example of **testing a connector**:


```
curl -X 'POST' \
  '<https://<APPLIANCE_ADDRESS>/v3/connectors/2-DATABASE-23/test>' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>'
```

9.4.3 Execute job

 Compliance Engines limit any syncing operations while a profiling or masking job is running. When using DCT job **move**, **execute**, or **migrate** operations, please ensure that the target Compliance Engine is in an idle state. Future Compliance enhancements to DCT will remove this limitation.

The **Execute endpoint** creates a duplicate of a job while maintaining a single reference for both job instances. This operation supports Horizontal Scale workflows, as DCT will maintain a singular reference for all instances of a job across any number of connected Compliance engines.

As part of this endpoint, DCT will maintain all of these job instances in sync, so they can all be controlled from a single configuration point (it's recommended to dedicate a select engine or set of engines to the creation and updating of masking jobs and dependencies) and any changes are automatically propagated to the other job instances at the time of the next job execute operation. This enables users to identify a masking job on a configuration engine, copy it over to a dedicated compute engine (or set of engines), and run that job at a regular cadence through DCT. Whenever the job needs to be updated, the user simply updates the job on the configuration engine.

 Since all jobs connected via the job execute operation are under a single reference, every time a job is run, its run statistics will report back to DCT and will be recorded under that singular job reference.

Executing a MaskingJob requires only a reference to a target engine as input. DCT will take care of syncing the job to the target engine and executing it. DCT will create and manage the environment where the job is copied onto.

Example of **executing a MaskingJob on engine with ID 2**:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead/execute' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
  "engine_id": "2"
}'
```

This will return a DCT job that can be further polled for status updates. The job will only transition to the COMPLETED state when the entirety of the sync and execution has completed on the target engine.

When a MaskingJob is executed via DCT and the job is synced to the target engine, the default Connector is used for execution. Masking job sync never copies credentials, for security reasons. Since having credentials set on the target connector is required for execution, DCT enables this by allowing users to store connector credentials within DCT itself. A DCT MaskingJob now contains properties for the connector credentials. The expectation is that users will pre-store the credentials by using the UPDATE API on the MaskingJob. MaskingJob execution has a hard requirement that credentials be saved within a MaskingJob prior to allowing execution.

Example of **updating a MaskingJob with connector credentials**:

```
curl -X 'PATCH' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "connector_username": "user123",
    "connector_password": "password123"
  }'
```

Once a MaskingJob execution has been initiated, the EXECUTION APIs can be used to view and cancel running executions as well as search through execution history. Note that canceling an execution is a best-effort action that does not interrupt any of the job sync that may occur prior to the execution.

Example of **searching for executions of a particular MaskingJob**:

```
'https://<APPLIANCE_ADDRESS>/v3/executions/search' \
-H 'accept: application/json' \
-H 'Authorization: <API_KEY>' \
-H 'Content-Type: application/json' \
-d '{
  "filter_expression": "masking_job_id eq '\''d53812ce-9186-485d-
a388-44bc52087ead'\''"
}'
```

Example of **canceling an execution if and only if it is in the RUNNING state (denoted by the expected_status query parameter)**:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/executions/11397caa-6006-4eba-b575-ae3ad00c3762/cancel' \
  -H 'accept: */*' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
    "expected_status": "QUEUED"
  }'
```

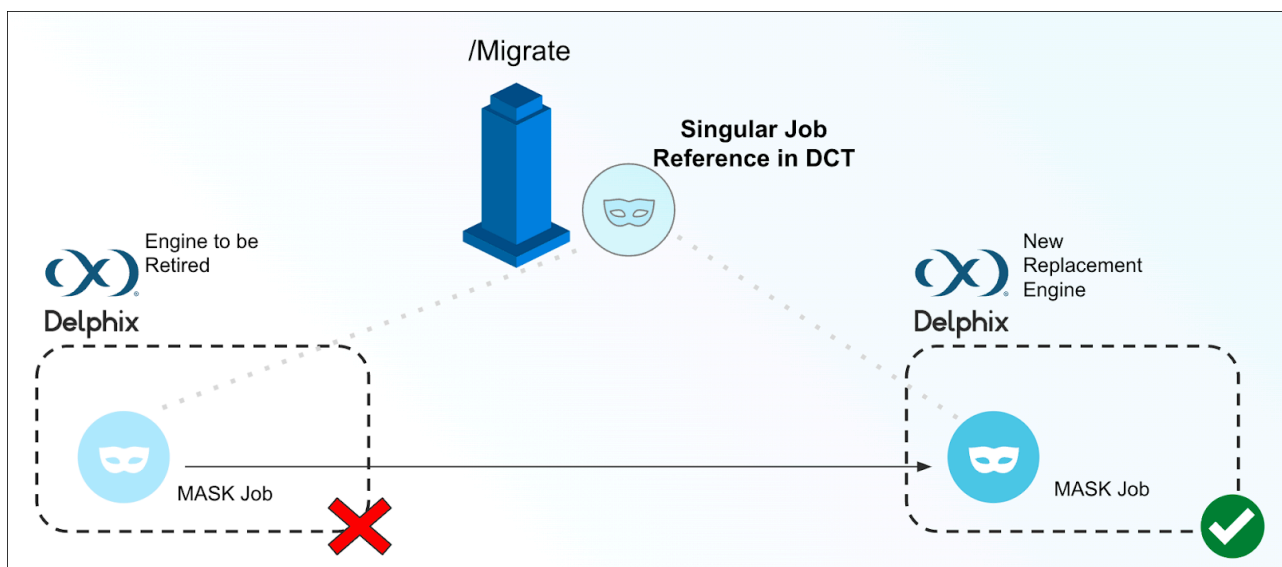
9.4.3.1 User interface

Execute a compliance job with the **Execute** action, available in the action menu on the top right corner of the job details page. This will open a window that lists Compliance engines with which a job needs to be executed. Once selected, click the "Execute" button to start the job on the selected engine. The screenshot below shows a selected engine.

9.4.4 Migrate job

⚠ Compliance Engines limit any syncing operations while a profiling or masking job is running. When using DCT job **move**, **execute**, or **migrate** operations, please ensure that the target Compliance Engine is in an idle state. Future Compliance enhancements to DCT will remove this limitation.

The **Migrate endpoint** moves a job from one engine to another without any duplicates. This endpoint is useful for consolidating masking jobs (i.e. moving jobs to a fresh engine ahead of the original being retired or consolidating two development engines into a single one for administrative simplicity). This means that a job will continue to have only a single instance with no additional jobs being created. This job will maintain its same reference within DCT.



Example of **finding all MaskingJobs originating from engine with ID 2:**

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/source-engines/search' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
```

```
"filter_expression": "source_ignite_id eq '\''2'\''"
}'
```

Example of **migrating a MaskingJob to new source engine with ID 3 and placing it in the 'prod-env' environment**:

```
curl -X 'POST' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead/migrate' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>' \
  -H 'Content-Type: application/json' \
  -d '{
  "target_engine_id": "3",
  "target_environment_id": "prod-env"
}'
```

9.4.5 Delete job

Calling the DELETE API on a MaskingJob will effectively remove the record from DCT (and its execution history) as well as delete the actual masking job on the source engine and on any other engine where the job has been copied to (as a result of execution). The API includes a force option to prevent the action from failing in the event that an engine is unreachable.

Example of **deleting a MaskingJob with the force option**:

```
curl -X 'DELETE' \
  'https://<APPLIANCE_ADDRESS>/v3/masking-jobs/d53812ce-9186-485d-a388-44bc52087ead?force=true' \
  -H 'accept: application/json' \
  -H 'Authorization: <API_KEY>'
```

This will return a DCT Job that can be further polled for status updates. Note that if the force option is used and there are ignored errors, details about those errors will be included in the **error_details** and **warning_message** fields of the DCT Job as follows:

```
{
  "job": {
    "id": "722ba51cf70e4e32adbd192b07304bb5",
    "status": "COMPLETED",
    "type": "MASKING_JOB_DELETE",
    "error_details": "Unable to connect to the engine.",
    "warning_message": "Failed to remove local MaskingJob, engineId: 3
localMaskingJobId: 7.",
    "target_id": "d53812ce-9186-485d-a388-44bc52087ead",
    "start_time": "2022-01-02T05:11:24.148000+00:00",
    "update_time": "2022-01-02T06:11:24.148000+00:00"
  }
}
```

9.5 Read-only algorithms

9.5.1 Overview

The DCT 13.0.0 release introduces a new functionality that allows users to explore the masking algorithms supported by DCT, as well as algorithms imported from a Continuous Compliance Engine.

DCT is equipped with a collection of plugins that support masking algorithms, each containing multiple built-in algorithms by default. In DCT version 13.0.0, the shipped algorithm plugin versions range from 1.8.0 to 1.21.0 and from 17.0.0.0 to 19.0.0.0.

In DCT, each algorithm may undergo multiple revisions, distinguished by its configuration and algorithm key. Connecting a Continuous Compliance Engine to DCT results in the importation of both built-in and custom algorithms from the engine. If an algorithm already exists in DCT, new revisions are created under the algorithm to accommodate those imported from the engine.

9.5.1.1 Feature Limitations

In DCT 13.0.0, all algorithms and their revisions are **read-only**. Users cannot create, delete, or modify algorithms or revisions. Users can, however, add or remove tags to algorithms and revisions, in addition to changing the name and note of revisions.

9.5.1.2 Algorithms page

- **Algorithm list view**
 - Displays a list of available masking algorithms.
 - Each algorithm entry shows:
 - Algorithm name.
 - Framework name.
 - Description.

- Tags.
- Clicking an algorithm or the **View** button redirects to its details overview page.
- **Filter and search functionality**
 - Allows users to filter and search through the list of algorithms based on names, plugin, framework, or tags.

The screenshot displays the 'Algorithms' page in the Data Control Tower interface. The page title is 'Algorithms' and the subtitle is 'View and manage algorithms used to mask sensitive data.' A search bar is located at the top right of the main content area. Below the search bar is a table listing various algorithms. Each row in the table includes the algorithm's name, the plugin it uses, the framework, a brief description, and a 'Tags' button. A 'View' button is also present at the end of each row. The table contains the following data:

Name	Plugin	Framework	Description	Tags
CommentLookup	dipx-core	Secure Lookup	Contains 1 generic value	Add Tags > View >
dipx-core:CM Alpha-Nu...	dipx-core	Character Mapping	An instance of the chara...	Add Tags > View >
DummyHospitalNameL...	dipx-core	Secure Lookup	Dummy (non real) hospi...	Add Tags > View >
SchoolNameLookup	dipx-core	Secure Lookup	School and Universities ...	Add Tags > View >
dipx-core:CM Numeric	dipx-core	CM Numeric	A character mapping (C...	Add Tags > View >
LastCommaFirstLookup	dipx-core	Secure Lookup	Last name then a comm...	Add Tags > View >
RepeatFirstDigit	dipx-core	Repeat First Digit	This algorithm masks a ...	Add Tags > View >
dipx-core:Phone US	dipx-core	Phone	The Phone US algorithm...	Add Tags > View >

At the bottom of the page, there is a pagination indicator showing '1 to 50 of 53' and 'Page 1 of 2'.

9.5.1.3 Algorithm Details Overview page

- **Algorithm Overview section**
 - Displays the selected algorithm's name and plugin.
 - Shows a brief description of the algorithm and its purpose.
- **Revisions list**
 - Lists all revisions of the selected algorithm.
 - Each revision entry includes the revision name.
- **Action buttons** (if applicable)
 - Buttons for adding/removing tags.
 - Options to change the name and note of revisions (reflecting the read-only nature of the algorithms).

Algorithms > dlpx-core:CM Digits

< Back **dlpx-core:CM Digits**
7cce277f-d7fb-4e82-92ad-7eca4b31658b

Overview Revisions

Plugin **dlpx-core**

Framework **Character Mapping**

Details


Description
An instance of the character mapping (CM) framework covering numeric digits [0-9]. This algorithm is appropriate for masking text data types when it is desirable to preserve the original non-digit content of each value. It produces a new, stable, and collision-free masked result for each input when applied to text data-types, but may produce collisions when applied to numeric data types.

Tags [Edit](#)

No tags added

9.5.1.4 Algorithm Revisions page

- **Revision list:**
 - Shows a list of the Revisions, along with columns for name, note, engine (which engine this revision comes from), and tags.
- **Import information:**
 - If the revision is imported from a Continuous Compliance Engine, additional details are shown, like the source of import and any custom modifications.
- **Tag management:**
 - Options to add or remove tags from the revision.
 - Reflects the user's ability to manage tags despite the read-only status of algorithms.
- **Name and note editing:**
 - Fields allowing users to change the name and add notes to the revision.



DATA CONTROL TOWER

Home Data **Compliance** Insights Admin

Algorithms > dlp-core:CM Digits

< Back **dlpx-core:CM Digits**

92bb218b-c730-4dec-a03a-5d0bf24f517a

Overview **Revisions**

Search Search (All) ▾

Name	Note	Engine	Tags
CM Digits-dlp-core-19.0.0		masking1	Add Tags > Actions ⋮
CM Digits-dlp-core-18.0.0		masking-18001	Add Tags > Actions ⋮
CM Digits-dlp-core-1.15.0		masking-9001	Add Tags > Actions ⋮

1 to 3 of 3 < Page 1 of 1 >

10 Hyperscale Orchestrator UI

10.1 Overview

The Hyperscale UI is *Controlled Availability*, meaning that Delphix highly encourages interested customers to reach out to Delphix Product and Engineering for a guided onboarding. Please contact your account team if interested.

To disambiguate object nuances between Compliance Engines, Hyperscale Orchestrators, and DCT, please refer to the following definitions:

- **Hyperscale dataset**
 - A Hyperscale dataset is a granular unit that defines source and target connector information alongside a defined database schema to be masked (the inventory on a Compliance Engine).
- **Hyperscale job**
 - A Hyperscale-specific object that is comprised of a rule set along with additional running parameters such as engines to be used and memory settings.
- **Hyperscale job execution**
 - The execution information such as duration, rows masked, etc. that is sent back to DCT for reporting.

Data Control Tower has the ability to connect to one or more [Hyperscale Compliance Orchestrators](#)⁸⁷, to serve as a complimentary user interface. The current Compliance Engine/Hyperscale Orchestrator/DCT workflow is as follows:

1. Build the foundational objects (rule sets, connectors, inventories, etc.) within the Compliance Engine.
2. Build a data set on the Hyperscale cluster by either:
 - a. Importing a compliance job from a Compliance Engine.
 - b. Building the dataset via the connector, dataset, and jobs API on the Hyperscale cluster (DCT will read the new Hyperscale jobs and display them.)
3. Set up the DCT-to-Hyperscale connection by following simple steps in the [Connecting/authenticating](#)⁸⁸ page. DCT will then create references for any data sets (registered as Hyperscale jobs in DCT), cluster details (e.g. mount points, etc.), and pre-existing engine connections to the Hyperscale Orchestrator.
4. Create a Hyperscale job on DCT by selecting an existing masking job and importing the job's inventory and connector information into DCT to represent a hyperscale job. Users have the option to further refine the job details such as defining table-specific masking parameters. A job will appear as a special Hyperscale job on the compliance jobs page.
5. Execute a Hyperscale job by selecting the **run** option on the Hyperscale jobs detailed view.

⁸⁷ <https://hyperscalemasking.delphix.com/docs/latest/>

⁸⁸ <https://dct.delphix.com/docs/10.0.0/connecting-authenticating>

- DCT will recognize this job execution and will present the execution details under the specific Hyperscale Orchestrator’s **Details** page.

10.2 Implementation introduction

The first step in working with Hyperscale is to register a Hyperscale Orchestrator with DCT. DCT connects to all Hyperscale Orchestrators over HTTPS, some configurations might be required to ensure DCT can communicate successfully.

Begin the Hyperscale Orchestrator registration flow, as well as view currently registered Hyperscale Orchestrators under the **Compliance** -> **Hyperscale Orchestrators** section. The registration wizard will guide you through the steps of the process, some of which are described in further detail below.

10.2.1 Truststore for HTTPS

If the CA certificate that signed the Hyperscale Orchestrator’s HTTPS certificate is not a trusted root CA certificate present in the JDK, then custom CA certificates can be provided to DCT. If these certificates are not provided, a secure HTTPS connection cannot be established and registering the Hyperscale Orchestrator will fail. The `insecure_ssl` Hyperscale Orchestrator registration parameter available as a checkbox option in the registration wizard can be used to bypass the check, however, this should not be used unless the risks are understood.

Get the public certificate of the CA that signed the Hyperscale Orchestrator’s HTTPS certificate in PEM format. Your IT team might be required to get the correct certificates. Base64 encode the certificate with:

```
cat mycertfile.pem | base64 -w 0
```

Copy the Base64 encoded value from the previous step and configure in the `values.yaml` file under the `truststoreCertificates` section. For example, the section might look like this:

```
truststoreCertificates: <certificate_name>.cert: <base64 encode certificate string
value in single line>
```

`<certificate_name>` can be any logically valid string value, like “hyperscale”.

All certificates configured in the `truststoreCertificates` section will be read and included in the `trustStore`, which would then be used for SSL/TLS communication between DCT and Hyperscale.

10.2.2 Authentication with Hyperscale Orchestrators

All authentication with a Hyperscale Orchestrator is done with an API Key corresponding to an admin Hyperscale user. The API will be stored and encrypted on DCT itself. The use of a vault to store Hyperscale credentials is currently not supported.

10.2.3 Hyperscale deployment type

A deployment type designation must be made for a Hyperscale Orchestrator at the time of registration. The registration UI wizard will present the data type options that are supported. This type must be in line with the type of data source the Hyperscale Orchestrator has been set up to mask (eg, Oracle).

10.2.4 Editing and unregistering Hyperscale Orchestrators

Making edits to a Hyperscale Orchestrator's configuration can be done via the UI, in the Orchestrator's overview page via the Edit action in the Details tile.


The screenshot shows the 'Update Hyperscale Orchestrator Configuration' dialog box. The background shows the 'Details' section of a Hyperscale Orchestrator page with the following information:

- Created:** Jan 17, 2024 8:15 AM PST
- Data Type:** Oracle
- Host:**
- Insecure SSL:** True
- Unsafe Hostname Check:** True
- Truststore Filename:** None

The dialog box contains the following fields and options:

- Host/IP Address (text input)
- API Key (password input, masked with asterisks)
- Data Type (dropdown menu, currently set to Oracle)
- Insecure SSL
- Unsafe SSL Hostname Check
- Truststore Filename (text input)
- Truststore Password (password input)
- Buttons: Cancel, Update


Unregistering a Hyperscale Orchestrator and removing all its data from DCT can be done via the `DELETE` API or directly in the UI via the Actions menu.



10.3 Managing Hyperscale objects

10.3.1 Introduction

After registering a Hyperscale Orchestrator, DCT will begin to ingest all relevant data discovered on that Hyperscale instance. This data, in turn, becomes modeled in the DCT world with APIs and a UI to (and manage to a limited extent). Currently, a lot of this data can be navigated to and viewed from the main overview details page of a Hyperscale Orchestrator.



10.3.2 Hyperscale executions

A view of all current and past Hyperscale job executions present on a particular Hyperscale Orchestrator can be found under the Job Executions tab. The **View** link will display additional information, especially for in-progress or failed executions.

Hyperscale Orchestrators > hyperscale-prod
hyperscale-prod
ID: b5d7299c-a8b1-468c-ba94-9c88e466b90c

Overview Job Executions Associated Jobs Engine Pool Mount Points Connector Configurations Access

Search Search (All) ▾

Status	Job Name	Start Time	End Time ↑	Duration	Rows Masked	
Running	the-big-job	Aug 21, 2023 12:40 ...	-	18 days 7 hrs 31 mins	201	View >
Completed	the-job	Aug 8, 2023 2:43 A...	Aug 8, 2023 2:48 A...	4 mins 51 secs	2	View >
Failed	the-big-job	Aug 18, 2023 1:28 A...	Aug 18, 2023 1:29 A...	52 secs		View >
Failed	the-big-job	Aug 18, 2023 1:53 A...	Aug 18, 2023 7:35 A...	5 hrs 41 mins 56 secs	1999	View >
Failed	the-big-job	Aug 21, 2023 12:22 ...	Aug 21, 2023 12:25 ...	2 mins 57 secs	451	View >

1 to 5 of 5 < Page 1 of 1 >

10.3.3 Hyperscale jobs

A view of all Hyperscale jobs discovered on a Hyperscale Orchestrator can be found under the **Associated Jobs** tab.

DATA CONTROL TOWER

[Home](#)
[Data](#)
[Compliance](#)
[Insights](#)
[Admin](#)

Hyperscale Orchestrators > hyperscale-prod

hyperscale-prod

ID: b5d7299c-a8b1-468c-ba94-9c88e466b90c

Overview
Job Executions
Associated Jobs
Engine Pool
Mount Points
Connector Configurations
Access

Search (All)

▼

Last Run Status	Name	Last Completed	Tags
Running	the-big-job	None	Add Tags > View >
Completed	the-job	Aug 8, 2023 2:48 AM PDT	Add Tags > View >

1 to 2 of 2
< Page 1 of 1 >

The **View** link will take you to the **Compliance Jobs** details page where more information about the job can be found.

DATA CONTROL TOWER

[Home](#)
[Data](#)
[Compliance](#)
[Insights](#)
[Admin](#)

Compliance Jobs > the-big-job

the-big-job

b5d7299c-a8b1-468c-ba94-9c88e466b90c-2

Overview
History
Access

In Progress

Started at 12:40 AM

Run Time: 18 days 7 hrs 39 mins

Progress: 31%

31%

[View Job Execution Details >](#)

Details

Type

Hyperscale

Description

Job created by Hyperscale Masking

Data Retention Setting

No

Hyperscale Data Set Details

Mount Point

test_mount

Tags Edit

No tags added

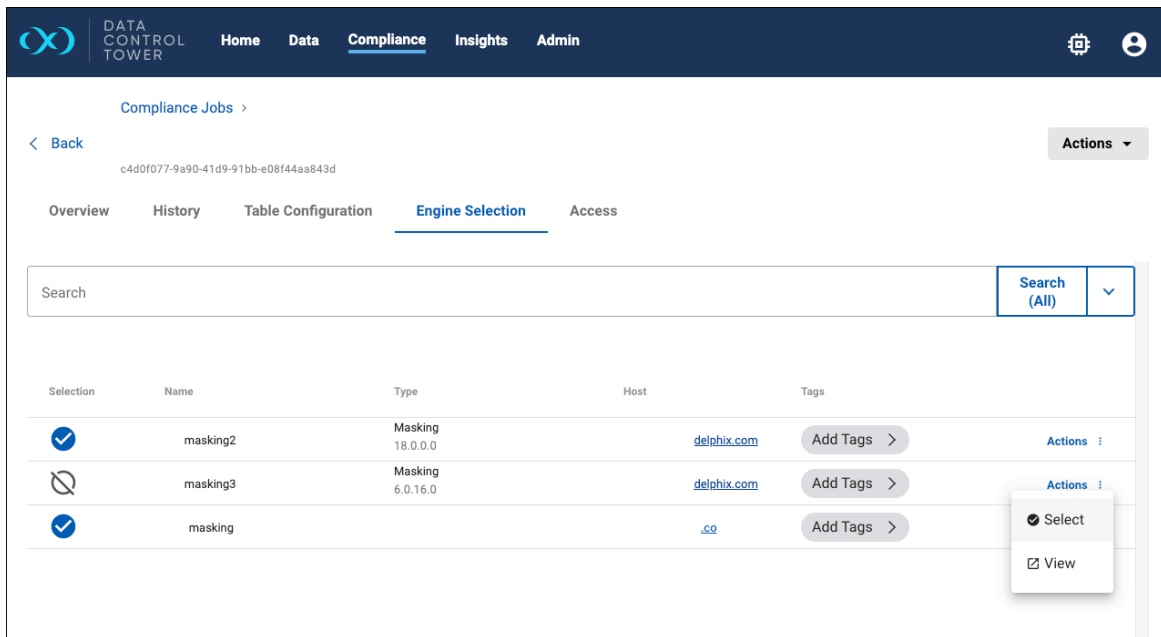
- Hyperscale jobs are not separate entities in DCT, but rather are combined with standard Compliance jobs. Hence, this Hyperscale job details page can be navigated to from **Compliance** -> **Compliance Jobs**. There is a `type` attribute (Hyperscale or standard) which helps differentiate between the different types.

Configuration and dataset details can be edited directly in the job's Overview page.

The screenshot displays the 'Update Configuration' modal window in the Data Control Tower. The modal contains a text input field for a description, currently containing 'Job created by Hyperscale Masking'. Below the input field are three radio button options under the heading 'Retain Execution Files': 'Always', 'Never' (which is selected), and 'On Errors'. At the bottom right of the modal are 'Cancel' and 'Update' buttons. The background shows the job overview page for a job named 'the-big-job' with a progress bar at 31% and various configuration details like 'Type: Hyperscale', 'Description: Job created by Hyperscale Masking', and 'Mount Point: test_mount'.

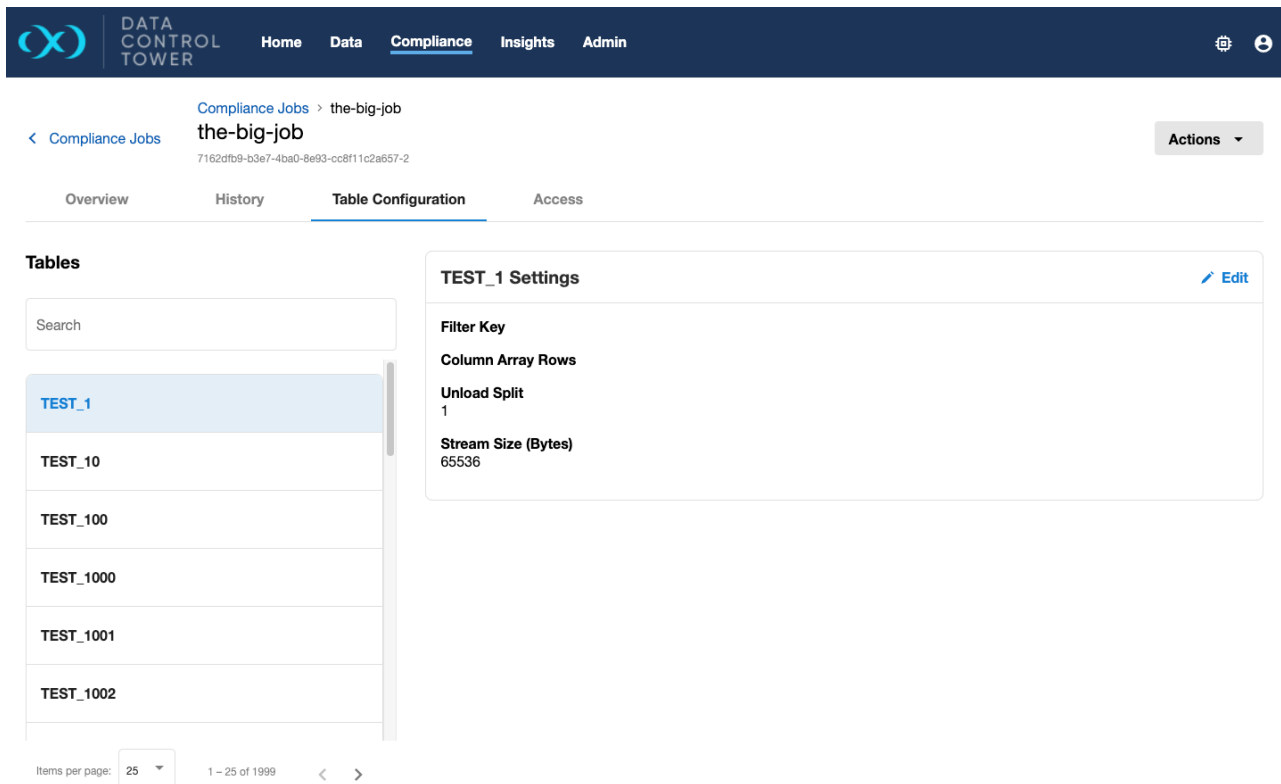
10.3.4 Hyperscale job engine selection

The Engine Selection tab allows you to manage the compliance engines that the job can use for masking. The table will list **all** compliance engines that are part of the Hyperscale Orchestrator's engine pool. Each compliance engine will either be Selected or Not Selected for the job. The state of each engine is indicated by the Selection column and can be changed via the Actions menu.




10.3.5 Hyperscale job table configuration

The Table Configuration tab contains the early stages of inventory management. You can see the Hyperscale dataset tables and their settings in the screenshot below. The left side shows the list of tables and allows you to search and paginate through the results.




The right side shows you the settings and allows you to make changes via the **Edit** action.



10.3.6 Executing Hyperscale jobs


A Hyperscale job can be executed via the Actions menu.



Once a job has started, a progress bar with details will appear in the Overview page.

Click the **View Job Execution Details** link to go to the Execution Details page to see job execution progress, as well as an Actions menu to **Stop** or **Re-Run** a job.

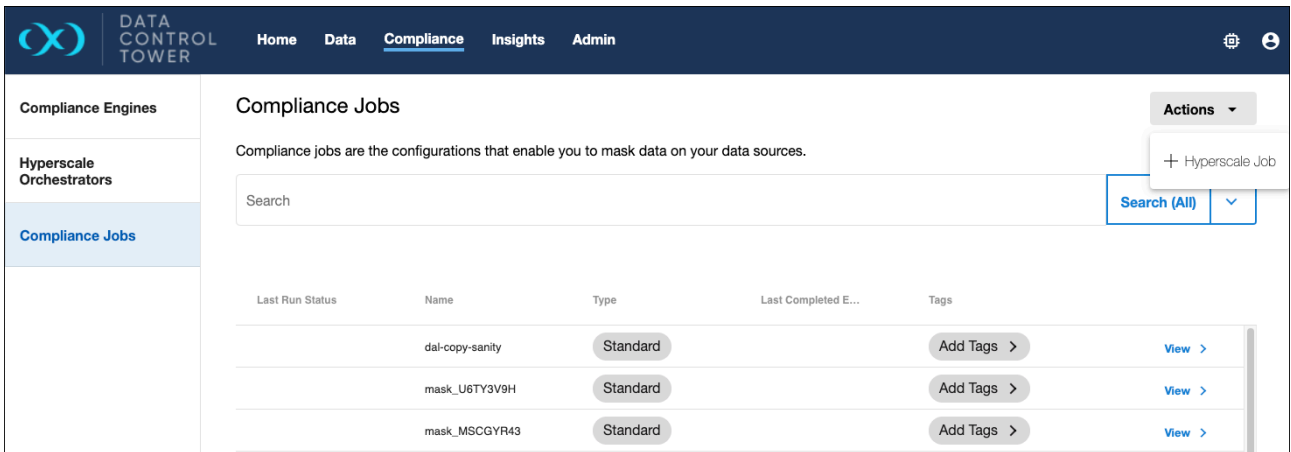
Details of a failed job will also be displayed in the Execution Details page. If the job was configured to retain execution data, the Clean option in the **Actions** menu can be used to discard this data.



10.3.7 Creating Hyperscale jobs

New Hyperscale jobs can be created via DCT when you have a database compliance job setup on a Compliance Engine and need to use the same masking inventory in a Hyperscale job. DCT will export the masking job details from the Compliance Engine and import them into the Hyperscale Orchestrator. The result is a new set of Hyperscale connectors, a dataset, and a Hyperscale job.

The **+ Hyperscale Job** button in the Actions menu will initiate a wizard that walks you through the creation process.



After setting the new job name, description, and tags, the wizard will prompt you to select a **source** compliance job. This compliance job must be the source job on your compliance engine whose inventory you want to import. You must then explicitly select a Hyperscale Orchestrator to create the job on, along with the **MountPoint** to use. Finally, set any and all configuration settings that are relevant to the job you want to create.

Create Hyperscale Job ✕

Details

Configuration

Summary

Configuration

Select Source Compliance Job

mask_V0203V31 Select

Source Credentials

delphix.com

Source Username

Source Password

Target Credentials

delphix.com

Target Username

Target Password

Select Hyperscale Orchestrator

em-hyperscale Select

Select Mount Point Infrastructure

test_mount Select

Default Stream Size (Bytes)

The stream size to be used by the sqldr Oracle utility which specifies the size (in bytes) of the data stream sent from the client to the server.

Default Unload Split

The number of unloaded files to be generated from the source database.

Retain Execution Files

Always

Never

On Error

Advanced Configuration ▼

Cancel
Back
Next
Submit

By default, DCT will create the Hyperscale job using all available Compliance Engines in the orchestrator's pool. If changing this selection is required, it can be done after job creation in the Engine Selection tab of the job's details page.

Currently DCT does not validate or enforce required input parameters. The requirements to create a Hyperscale job will differ depending on the Hyperscale Orchestrator version and database type.

As a reminder, the import process will not include any secure credentials for connectors. You have the option to set the connector credentials up front when creating the job (see screenshot above). Otherwise, you must independently find the imported connectors and explicitly set the credentials after the job has been created.

If your source compliance job is using any global objects such as Algorithms, those global objects must be pre-synced to all engines in the orchestrator's engine pool. This can be done via the **Manually Deploy Job Dependencies** option in the Hyperscale Orchestrator's **Action** menu.

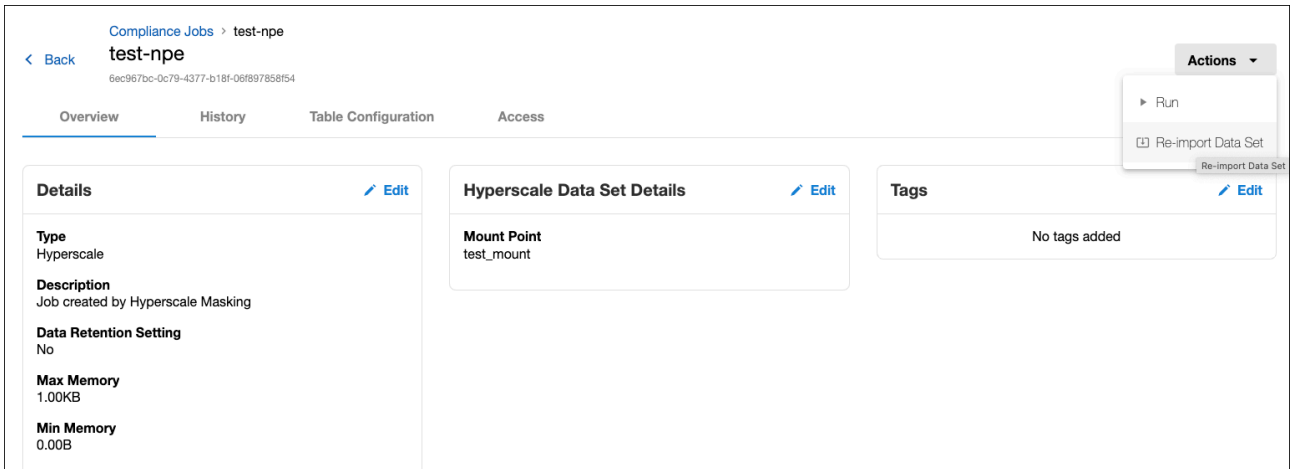
The wizard will walk you through the necessary steps which includes selecting the source compliance engine along with one or more target engines.

Name	Engine Version	Host Name	Tags
<input type="checkbox"/> eyal-masking		eyal-masking2.dlpxdc.co	Add Tags >
<input checked="" type="checkbox"/> em-masking2	18.0.0.0	em-masking2.dcol2.delphix...	Add Tags >
<input type="checkbox"/> em-masking3	6.0.16.0	em-masking3.dcol2.delphix...	Add Tags >
<input type="checkbox"/> UI Masking	20.0.0.0	ui-masking.dcol2.delphix.com	Add Tags >

Please refer to the [Hyperscale product documentation](https://documentation.delphix.com/docs/hyperscale-compliance)⁸⁹ for the Hyperscale Orchestrator version you are running to learn all about the specifics of creating jobs in Hyperscale.

⁸⁹ <https://documentation.delphix.com/docs/hyperscale-compliance>


Updates to the original source compliance job can be done anytime after initial import and a re-import can be done via DCT. This will update the Hyperscale job’s existing dataset with the refreshed ruleset from the source compliance job.



Note that this action will keep existing table configurations intact and the initial default settings for **unload split** and **stream size** will apply only to new tables. These settings can be changed for any table afterwards in the job’s **Table Configuration** details.

10.3.8 Hyperscale Compliance Engines

A view of all Compliance Engines registered with a Hyperscale Orchestrator can be found under the **Engine Pool** tab.



When DCT discovers the engines registered with a Hyperscale Orchestrator, it will create DCT `RegisteredEngine` entities out of them (if they do not already exist in DCT, as uniquely identified by the hostname). The result is a unified model where DCT `RegisteredEngine` objects are what make up the

engine pool in a Hyperscale Orchestrator. The same engines will appear under the **Compliance -> Compliance Engines** page.

- Engine credentials will not be retrieved from the Hyperscale Orchestrator, so the created engines must be updated with credentials, along with any relevant security settings. The discovered engines will remain in the OFFLINE status until updated.

Compliance Engines Connect Engine

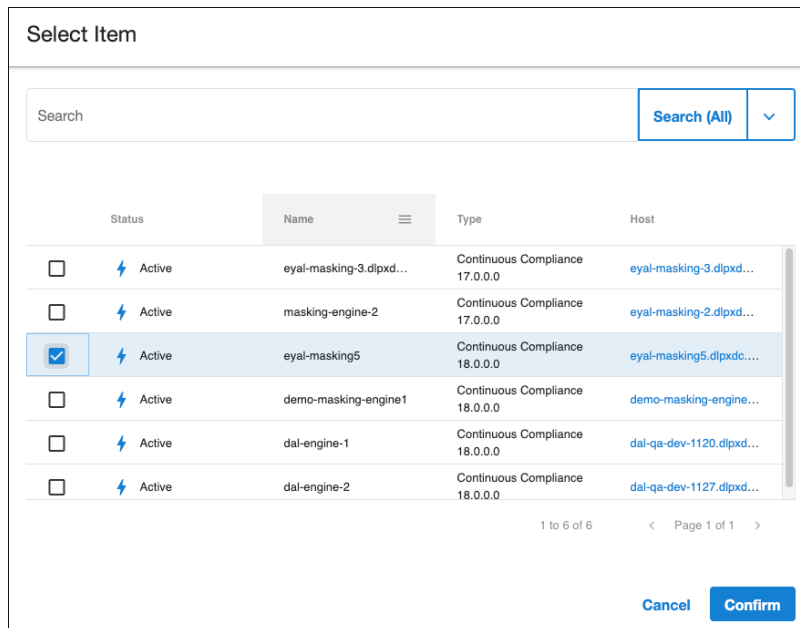
Overview of all compliance engines in your Delphix ecosystem.

Search Search (All) ▾

Status	Name/UUID ↑	Type	Running Jobs	Masking Usage/Total	Masking Available Cores	Tags
● Online	em-masking	Masking 15.0.0.0	0	0% 2.00GB	2	Add Tags > View >
● Offline	eyal-masking	Masking 15.0.0.0	0	0% 2.00GB	2	Add Tags > View >


- The created DCT engine entities are permanent, in the sense that even if the origin Hyperscale Orchestrator is unregistered, the engines in DCT remain as if they had been registered independently of Hyperscale.

Adding new engines to the engine pool can be done by clicking the **+ Engine** button. Only Compliance Engines that have already been registered with DCT can be added to a Hyperscale Orchestrator's engine pool:



Note that this operation will, in turn, register the Compliance Engine with the Hyperscale Orchestrator. The Compliance Engine name, credentials, and configuration settings will be set according to DCT's record.

Removing an engine from the engine pool can be initiated with the **Remove** button via the Actions menu:



This will completely unregister the Compliance Engine from the Hyperscale Orchestrator.

10.3.9 Hyperscale mount points

A view of all mount points on a Hyperscale Orchestrator can be found under the **Mount Points** tab.

The screenshot shows the 'Mount Points' tab for a Hyperscale Orchestrator named 'hps-mssql'. The interface includes a search bar, a '+ Mount Point' button, and a table of existing mount points.

Mount Point Name	Host Address	Type	Path	Actions
my-mp1-1	mp1.delphix.com	NFS3	/tmp/foo	Actions
test_mount	e-nfs.dlpxdc.co	NFS4	/var/tmp/masking-mount/	Actions
test_mount_2	e-nfs3.dlpxdc.co	NFS3	/var/tmp/masking-mount/	Actions

Page 1 of 1

Creating new mount points can be done via the **+ Mount Point** button.

The screenshot shows the 'Add Mount Point' dialog box. It contains the following fields and options:

- Mount Point Name (text input)
- Host Address (text input)
- Path (text input)
- Type (radio buttons):
 - CIFS
 - NFS3
 - NFS4
- Options (text input)
- Buttons: Cancel, Add Mount Point

This will result in a new mount point being created directly on the Hyperscale Orchestrator.

Editing and **Deleting** a mount point can be done for a particular row via the Action menu in the last column of the table.

The screenshot shows the 'Mount Points' tab for the 'hps-mssql' orchestrator. The interface includes a search bar, a '+ Mount Point' button, and a table with the following data:

Mount Point Name	Host Address	Type	Path
my-mp1-1	mp1.delphix.com	NFS3	/tmp/foo
test_mount	e-nfs.dlpxdc.co	NFS4	/var/tmp/masking-mount/
test_mount_2	e-nfs3.dlpxdc.co	NFS3	/var/tmp/masking-mount/

Each row in the table has an 'Actions' menu with 'Edit' and 'Delete' options. The bottom of the page shows pagination: '1 to 3 of 3' and 'Page 1 of 1'.

10.3.10 Hyperscale connector configurations


A view of all connectors on a Hyperscale Orchestrator can be found under the **Connector Configurations** tab.

The screenshot shows the 'Connector Configurations' tab for the 'hps-mssql' orchestrator. The interface includes a search bar, a '+ Connector' button, and a table with the following data:


Connector ID	Source JDBC URL	Target JDBC URL
380b1797-89a0-4d75-a1b4-607935a31997-1	jdbc:sqlserver://10.110.229.231:49854;database=F...	jdbc:sqlserver://10.110.192.63:49854;database=Fe...
380b1797-89a0-4d75-a1b4-607935a31997-2	jdbc:sqlserver://10.110.229.27:59871;database=S...	jdbc:sqlserver://10.110.193.10:59871;database=S...
380b1797-89a0-4d75-a1b4-607935a31997-3	jdbc:sqlserver://10.110.242.215:59871;database=S...	jdbc:sqlserver://10.110.244.26:59871;database=S...
380b1797-89a0-4d75-a1b4-607935a31997-4	jdbc:sqlserver://10.110.249.67:59871;database=S...	jdbc:sqlserver://10.110.238.133:59871;database=S...

Each row in the table has an 'Actions' menu. The bottom of the page shows pagination: '1 to 4 of 4' and 'Page 1 of 1'.

Creating new connectors can be done via the **+ Connector** button.



Editing and Deleting a connector can be done for a particular row via the Actions menu in the last column of the table.



11 Integrations

Data Control Tower provides a global integration layer for a connected Delphix ecosystem, whether that is a single or dozens of globally distributed engines, DCT drives a scalable approach to integrating Delphix into any custom script or automation toolchain.

Aside from the comprehensive API layer (see [API references \(see page 295\)](#) for more detail), DCT powers automation through Delphix-built and supported integrations with popular applications such as Terraform, ServiceNow, etc.

To see a current list of Delphix integrations, please visit [Delphix Integrations](#)⁹⁰ for more detail.

⁹⁰ <http://ecosystem.delphix.com>

12 DCT concepts

12.1 Introduction

Data Control Tower (DCT) provides new and novel approaches to general Delphix workflows, delivering a more streamlined developer experience. This article will introduce these concepts to Delphix and how they work with DCT.



For **VDB Provisioning**, the UI supports these data platforms:

- Oracle Single Instance Single Tenant
- Oracle Single Instance Multi Tenant (for Linked CDB only)
- MSSql Single Instance


For **Infrastructure Connection Wizard**, only UNIX standalone environments can be added via the UI.

12.2 Concepts

12.2.1 Virtual Database (VDB) groups

Virtual Database (VDB) groups are a new concept to Delphix, which enable the association of one or more VDBs as a single VDB group. This allows for bulk operations to be performed on the grouped VDBs, such as bookmark, provision, refresh, rewind, and others. This will assist in complex application testing scenarios (e.g. integration and functional testing) that require multiple data sources to properly complete testing.

With VDB groups, developers can now maintain data synchronicity between all grouped VDBs, which is particularly useful for complex timeflow operations. For example, updating VDBs to reflect a series of schema changes across data sources, or to reflect an interesting event in all grouped datasets. In order to maintain synchronicity among grouped datasets, timeflow operations (refresh, rewind, etc.) must use a bookmark reference.



In the above example, a VDB Group reference is created for three VDBs. At the end of the above timeline group, a developer decides to rollback those VDBs to a previous snapshot. By issuing a single command via the VDB groups endpoint, DCT will move all three back, ensuring that they all maintain referential synchronicity.

Bookmarks and VDB groups are loosely related; a VDB group can exist in the absence of any bookmarks, and a bookmark can exist without any VDB group. It is important to note that the bookmark represents data, while the VDB group represents the databases to make this data available.

ⓘ DCT will automatically stop an operation from executing if one or more objects are incompatible (e.g. provisioning a VDB group into a set of environments, where one of the VDBs is incompatible, such as an Oracle on Linux VDB provisioned onto a Windows environment).

VDB groups based operations will return a single job to monitor the overall status of the series of individual VDB operations. If one of those individual operations is unable to complete, DCT will report a “fail”, but any individual operations that are able to successfully complete will still do so.

12.2.2 Comparing Self-Service containers to VDB groups


As mentioned above, VDB groups are a crucial DCT concept that enable Self-Service functionality outside of the Self-Service application. Consider VDB groups acting similarly to Self-Service containers, in that it provides grouping and synchronization among VDBs, but VDB groups can provide a more flexible approach for users. Here are some additional points for example:

- The same VDB can be included in multiple VDB groups
- Including a VDB in a VDB group does not prevent operations on the VDB individually
- VDBs can be added to or removed from VDB groups
- VDB groups do not have their own timeline

12.2.3 Bookmarks

DCT Bookmarks are a new concept that represents a human-readable snapshot reference that is maintained within DCT. This is not to be confused with Self-Service bookmarks, maintained separately within the Self-Service application. With DCT Bookmarks, developers can now reference meaningful data (e.g. capturing a schema version reference to pair with an associated code version, capturing test failure data so that developers can reproduce the error in a developer environment, etc.) and use those references for any number of use-cases (e.g. versioning data as code, quickly provisioning a break/fix environment with relevant data, etc.). DCT Bookmarks are compatible with both VDBs and VDB groups, and can be used as a reference for common timeflow operations such as:

- Provisioning a VDB or VDB group from a bookmark
- Refreshing a VDB or VDB group to a bookmark
- Rewinding a VDB or VDB group to a bookmark


 DCT Bookmarks have associated retention policies, the default value is 30 days, but policies can be customized anywhere from a day to an infinite amount of time. Once the Bookmark expires, DCT will delete the bookmark.

Bookmarks are compatible with individual VDBs and VDB groups. Bookmark Sharing is only available for engines on version 6.0.13 and above.

DCT Bookmarks, when created, initiate a snapshot operation on each and every VDB in order to maintain synchronicity between each VDB. In that same vein, bookmark-based VDB group operations will have each VDB-specific sub-process run in parallel (as opposed to sequentially) to reduce drift between grouped VDBs.

12.2.4 Jobs

Jobs in DCT are the primary means of providing operation feedback (PENDING, STARTED, TIMEDOUT, RUNNING, CANCELED, FAILED, SUSPENDED, WAITING, COMPLETED, ABANDONED) for top-level operations that are run on DCT. Top-level operations represent the parent operation that may have one or more child-based jobs (e.g. refreshing a VDB group is the parent job to all of the individual refresh jobs for the grouped VDBs under the VDB group reference).

 Top-level jobs will report a “FAILED” status if one or more child jobs fail. For child jobs that can complete, DCT will continue to complete those jobs even if a parent job reports a failure.

12.2.5 Tags

DCT Tags enable a new business metadata layer for users and consumers to filter, sort, and identify common Delphix objects, to power any number of business-driven workflows. A tag is comprised of a (Key:Value) pair that associates business-level data (e.g. location, application, owner, etc.) with supported objects. DCT 2.0 and above support the following Tags:

- Continuous Data Engines
- Environments
- dSources
- VDBs

Developers and administrators add and remove tags using tag-specific object endpoints (e.g. `/vdb/{vdbId}/tags`) and can leverage tags as search criteria when using the object-specific search endpoints (e.g. using filtering language to narrow results).

Some sample tag-based use-cases include:

- Refreshing all the VDBs owned by a specific App Team using an “Application: Payment Processing” tag. This would be accomplished by querying “what VDBs have the (Application: Payment Processing) tag” and feeding those VDB IDs into the refresh endpoint.
- Driving accountability for VDB ownership by tagging primary and secondary owners for each VDB (e.g. (primary_owner: John Smith), (secondary_owner: Jane Brown)). That way, if a VDB is overdue for a refresh, tracking down an owner is a simple tag query.

[-] Tags are registered as an attribute that is specific to an object as opposed to a central tagging service. As a result, tag-based querying can only be done on a per-object type basis.

A supported object can contain any number of tags.

12.2.6 Tag-based filtering

All taggable objects support tag-based filtering for API queries that adhere to the search standards documented in [API References](#) (see page 295). A few examples of how tag-based filtering can be used are as follows:

List all VDBs of type 'Oracle', of which IP address contains the '10.1.100' string and which have been tagged with the 'team' tag, 'app-dev-1'.

```
database_type EQ 'Oracle' AND ip_address CONTAINS '10.1.100' and tags CONTAINS { key EQ 'team' AND value EQ 'app-dev-1' }
```

12.3 Nuances

12.3.1 Stateful APIs

All applicable DCT APIs are stateful so that running complex queries against a large Delphix deployment can be done rapidly and efficiently. DCT accomplishes this by periodically gathering and hosting telemetry-based Delphix metadata from each engine.

12.3.2 Local data availability

DCT currently relies on existing Continuous Data and Compliance constructs around data-environment-engine relationships. This means that DCT operations require VDBs to live on the engine where the parent dSource lives and so on.

12.3.3 Engine-to-DCT API mapping

Wherever possible, DCT has looked to provide an easier-to-consume developer experience. This means that in some cases, an API on DCT could have an identical API on an engine. However, there are many instances of providing a higher level abstraction for ease of consumption; one example is the data inventory APIs on DCT (sources, dSources, VDBs), which are a simplified representation of data represented by the source, sourceconfig, and repository endpoints on the local engine (source, dSource, and VDB detail are all combined under those three endpoints).

12.3.4 Local references to global UUIDs

In order to avoid collision of identically-named and referenced objects, DCT generates Universally Unique Identifiers (UUID) for all objects. For existing objects on engines like dSources and VDBs, DCT will concatenate the local engine reference with the engine UUID (e.g. 'Oracle-1' on engine '3cec810a-ee0f-11ec-8ea0-0242ac120002' will be represented as 'Oracle-1-3cec810a-ee0f-11ec-8ea0-0242ac120002' on DCT).

12.3.5 Environment representations

Environments within Delphix serve as a reference for the combination of a host and instance. This is coupled with the fact that environments can be leveraged by multiple engines at the same time and that engines often have a specific context to some of the elements that comprise an environment. For example, an environment could have both an Oracle and ASE instance installed and that Engine A leverages an Oracle-based workflow and Engine B leverages an ASE workflow. DCT will create two identifiers to represent the specific host and instance combinations. Thus, in DCT, Engine A will be connected to a different uniquely identified Environment than Engine B.

As mentioned earlier with Engine-to-DCT API mapping, DCT aims to simplify the user experience with Delphix APIs by combining different Continuous Data endpoints into a simplified DCT API. The Environment API does this by combining environment, repository, and host endpoints so that writing queries against Delphix data is

a much simpler process. One example would be identifying all environments that have a compatible Oracle home for provisioning:

```
repositories CONTAINS { database_type EQ 'Oracle' and allow_provisioning EQ true AND  
version CONTAINS '19.2.3'}
```

12.3.6 Supported data sources/configurations

DCT is compatible with all Delphix-supported data sources and configurations.

12.3.7 Process feedback

Whenever a DCT request completes, it will return a JOB ID as its response. This Job ID can be used in conjunction with the jobs endpoint to query the operation status.

13 DCT Toolkit

13.1 Introduction

DCT Toolkit (`dct-toolkit`) is the command line application specifically designed to interact with a particular DCT instance. It aims to simplify and ease out interacting with DCT instances while making the process intuitive, with minimal configuration and setup. `dct-toolkit` abstracts all of the API level implementation complexities and provides a user friendly LINUX-like interface, which can be run from any terminal.

13.2 Compatibility

`dct-toolkit` is designed to be compatible with any DCT product version, but is fully supported from DCT version 8.0.0 onwards. While `dct-toolkit` can be used in testing or development environments against any DCT versions, Delphix highly recommends at least DCT version 8.0.0 for using it in the production environment. Some of the value added features are only supported from DCT 8.0.0 version onwards.

`dct-toolkit` does not need to be upgraded with every new version of DCT. All of the new APIs that are part of the newer DCT releases will automatically be available in `dct-toolkit`. This is possible because `dct-toolkit` reads the api specification from the DCT instance (configurable) it is configured to connect to and dynamically generates the list of commands and their respective options.

13.3 New features

13.3.1 Version 1.2.0

- **`--json/-js` option**

`dct-toolkit` now supports a `--json/-js` option for all of the commands. If this option is specified, all properties from the DCT response are printed in JSON format. This option is mutually exclusive with the `--all-columns/-A` and `--columns/-c` options, but can be used along with the `--jsonpath/-jp` and `--expand/-ex` options.

- **`--csv/-cs` option**

`dct-toolkit` now supports a `--csv/-cs` option for all of the commands. If this option is specified, all properties from the DCT response are printed in CSV format. This option can be used along with the `--columns/-c` option only to output particular properties from the DCT response. This option can also be specified along with the `--jsonpath/-jp` and `--expand/-ex` options.

- This option is only applicable to list API responses where the response is printed in table format by default. For all other APIs that return a single object, this option is ignored.

- **`create_config` and `encrypt_api_key` commands are now interactive**

While providing `apiKey` value in `create_config` and `encrypt_api_key` command, there is a risk that plaintext `apiKey` can leak via shell history. To protect against this, both of these commands are now made interactive. Providing an `apiKey` interactively does not leak it via shell history. Following examples demonstrates providing `apiKey` interactively. Once following commands are executed, `dct-toolkit` will ask and wait for the `apiKey` to be provided by the user before executing the command.

```
./dct-toolkit encrypt_api_key dctUrl=<DCT instance URL> apiKey
```

```
./dct-toolkit create_config dctUrl=<DCT instance URL> apiKey
```

13.4 Installation and setup

13.4.1 Installation

`dct-toolkit` is currently supported on Linux, Windows, and macOS platforms. There are no special system requirements for running it.

To install `dct-toolkit`, download the executable binaries from the [Delphix download site](#)⁹¹ and extract it to the machine from which it needs to run. Once the executable binary is extracted, grant the execute permission to the binary file.


On **Linux** and **macOS** this can be done from Terminal with the following command:

```
chmod 777 ./dct-toolkit
```

For **Windows** this can be done from cmd with the following command:

```
CACLS files /e /p <username>:F
```

Where `<username>` is the user to whom the permission must be granted and `F` is the **Full Control** permission that needs to be granted.

-  The binary is not currently verified for macOS and will give a warning that the binary is untrusted. To get around this issue, you will need to open the binary and explicitly whitelist it. One easy way to achieve this is by pressing the ctrl key and clicking on the binary file, then select Open. On the resulting warning message box, select Open. Now the binary is allowed to be executed from the terminal.

⁹¹ [https://download.delphix.com/folder/4721/Delphix%20Product%20Releases/DCT/Toolkit%20\(CLI\)](https://download.delphix.com/folder/4721/Delphix%20Product%20Releases/DCT/Toolkit%20(CLI))

13.4.2 Setup

A plain text configuration file is needed to get started. The recommended way to create this configuration file is via the `create_config` command, available in `dct-toolkit`.

For creating a typical configuration with only the required properties, use the following command:

```
dct-toolkit create_config dctUrl=<URL of DCT instance> apiKey=<api key value>
```

This will create the configuration file in the default location under the user's home directory. The full path for the configuration file will be `<users home directory>/.dct-toolkit/dct-toolkit.properties`.

The `create_config` command takes in the options mentioned below:

Option name	Required/optional	Description
<code>dctUrl</code>	Required	HTTP/S URL of the DCT instance.
<code>apiKey</code>	Required	API key of the account used to connect to the DCT instance pointed by the <code>dctUrl</code> option.
<code>configFileOutputLocation</code>	Optional	Fully qualified path of the properties file. If this option is not specified, the properties file will be created in the <code>.dct-toolkit</code> folder under the user's home directory. If this option is specified, the <code>DCT_TOOLKIT_CONFIG_FILE</code> environment variable needs to be exported to point to the config file, in order for the config file to be effective.
<code>apiVersion</code>	Optional	API version to be used for DCT APIs. If this is not set, the latest API version from the DCT instance will be used.
<code>apiYamlLocation</code>	Optional	Location of the local <code>api.yaml</code> file. If this is not set, the latest <code>api.yaml</code> file will be downloaded from the DCT instance that is configured.
<code>--insecureSSL</code>	Optional	Whether to use insecure SSL connection to the DCT instance. By default, this will be set to false.

<code>--unsafeHostnameCheck</code>	Optional	Whether to disable hostname verifier checks for SSL connection to the DCT instance. By default, this will be set to false.
<code>sslCertificate</code>	Optional	PEM format SSL certificate path to be used for SSL connection to DCT instance.
<code>logLevel</code>	Optional	Log level to set. Can be one of OFF, ERROR, WARN, INFO, DEBUG, TRACE, ALL. By default, the log level will be INFO.
<code>logDir</code>	Optional	Directory where the logs should be written. By default, logs will be written to the logs directory created under the <code>.dct-toolkit</code> folder in the user's home directory.

All of these options and help can be requested with the following command:

```
./dct-toolkit create_config -h
```

All of the above properties can also be individually specified using environment variables. Environment variable mappings for each of the above properties are as follows:

Property name	Environment variable name
<code>dctUrl</code>	DCT_TOOLKIT_DCT_URL
<code>apiKey</code>	DCT_TOOLKIT_API_KEY
<code>apiVersion</code>	DCT_TOOLKIT_API_VERSION
<code>apiYamlLocation</code>	DCT_TOOLKIT_API_YAML_LOCATION
<code>--insecureSSL</code>	DCT_TOOLKIT_SSL_INSECURE
<code>--unsafeHostnameCheck</code>	DCT_TOOLKIT_SSL_UNSAFE_HOSTNAME_CHECK

<code>sslCertificate</code>	<code>DCT_TOOLKIT_SSL_CERT</code>
<code>logLevel</code>	<code>DCT_TOOLKIT_LOG_LEVEL</code>
<code>logDir</code>	<code>DCT_TOOLKIT_LOG_DIR</code>

Configuring non-sensitive properties in a properties file and sensitive properties (like `apiKey`) via an environment variable is acceptable. Properties set via environment variable will take precedence over the ones specified in the properties file.

Once the required properties are available via either the properties file or the environment variable, `dct-toolkit` is ready for use.

13.5 Usage guide

All of the DCT APIs are available as commands in `dct-toolkit`. To find the list of all commands which are available, use the following command:

```
./dct-toolkit -h
```

This will print the list of all commands grouped by the appropriate entity names.

13.5.1 Examples

- The list of available options for a particular command can be requested with the following command:

```
./dct-toolkit <command_name> -h
```

- There are currently many options for provisioning a VDB, some options are only valid for a particular DB type. For example, the `unique_name` option is only applicable to Oracle DBs. For better ease of use, `dct-toolkit` has subcommands created under the provisioning command, with options that are relevant to that particular subcommand. For example, these are the subcommands under the `provision_vdb_by_snapshot` command:

```
oracle
sybase
mssql
appdata
postgres
```

For provisioning an Oracle VDB, use the following command:


```
dct-toolkit provision_vdb_by_snapshot oracle source_data_id=<dsourceId> --
auto_select_repository unique_name=<uniqueName>
```

For provisioning a Sybase-ASE VDB, use the following command:

```
dct-toolkit provision_vdb_by_snapshot sybase source_data_id=<dsourceId> --
auto_select_repository --truncate_log_on_checkpoint
```

- For getting the list of registered engines, use the following command:

```
dct-toolkit get_registered_engines
```

The above command will only print important fields in the response.

For printing all of the fields in the response, use the following command:

```
-A/--all-columns
```

```
dct-toolkit get_registered_engines -A
```

To list only the selected columns that are useful to the user, use the following command:

```
--columns/-c
```

```
dct-toolkit get_registered_engines --columns=id,name,cpu_core_count...
```

For advanced use cases, dct-toolkit also provides an option to specify a [jsonpath](#)⁹² expression to extract required objects from the JSON response. This can be requested with the following command:

```
--jsonpath/-jp
```

```
dct-toolkit get_registered_engines --jsonpath=<json path expression starting
with $>
```

Note, these options are available for all of the commands available in dct-toolkit.

- For commands that use a request body, dct-toolkit provides `-body`, which can be used to specify a JSON body (instead of using individual options).

⁹² <https://github.com/json-path/JsonPath>

```
dct-toolkit token_info --body=<JSON body to POST to DCT>
```

Help regarding the exact structure of JSON request and response body can be requested for all of the commands with the following command:

```
-jh/--jsonhelp
```

```
dct-toolkit token_info -jh
```

- `dct-toolkit` provides an easy to use and understand format for specifying tags. For example, the following command adds two tags – `{key=purpose, value=testing}` and `{key=key1, value=value1}` :

```
dct-toolkit create_environment_tags environmentId=<envId> tags purpose=testing key1=value1
```

- Some of the DCT APIs trigger an asynchronous job and return a `jobId` in response. `dct-toolkit`, by default, will wait for the asynchronous job to be completed and will report the end state of the job.
 - There is an option to **not** wait for the completion of an asynchronous job. Use the following command with the `--no-wait` option:

```
dct-toolkit refresh_environment environmentId=<envId> --no-wait
```

- If the user intends on specifying the API key via an environment variable, `dct-toolkit` provides a command to get the encrypted version of the API key.

```
./dct-toolkit encrypt_api_key apiKey=<DCT api key> dctUrl=<DCT instance URL>
```

13.5.1.1 `--json/-js` and `--csv/-cs` option examples

- For the `--json/-js` option, to print the DCT response as it is received from DCT instance users can use following command:

```
./dct-toolkit get_environments --json
./dct-toolkit get_environment_by_id environmentId=<environment Id> --json
```

- To print the repositories object from environment object in JSON format users can use following command:

```
./dct-toolkit get_environments --json --expand=repositories
```

- For the `--csv/-cs` option, to print the DCT response in CSV format users can use the following command:

```
./dct-toolkit get_environments --csv
```

- To print the repositories object from the environment object in CSV format users can use the following command:

```
./dct-toolkit get_environments --csv --expand=repositories
```

13.6 API key encryption

13.6.1 Overview

dct-toolkit currently authenticates HTTP requests with an API key. The API key can be configured in plain text or specified via an environment variable.

13.6.2 Implementation

Starting from 1.2.0 version, **dct-toolkit** supports encrypting an API key via a separate command and also via the `create_config` command used to generate the configuration file. The `create_config` command now stores the API key in encrypted format using the AES/GCM algorithm, where the key and IV are generated from the following two inputs:

1. Local hostname of the machine from where the **dct-toolkit** command is run.
2. DCT URL used to connect to the DCT instance.

For users who want to supply the API key via an environment variable, **dct-toolkit** provides the `encrypt_api_key` command which takes in the API key and DCT URL as inputs and returns an encrypted API key, which can then be exported to an environment variable. As an example, please refer to the section below.

i The encryption key (not to be confused with the DCT API key) used for encryption is reproducible, hence, it is possible to encrypt and decrypt the DCT API key without storing the encryption key on the file system. Anyone with access to the logic to create an encryption key and knows the inputs for that logic will be able to generate the encryption key, and consequently, will be able to decrypt the DCT API key.

13.6.3 Backward compatibility

The `dct-toolkit` API key encryption feature is fully backward compatible with previous versions of `dct-toolkit`. `dct-toolkit` first tries to decrypt the API key; if it fails to decrypt the API key, then the key is used as configured.

13.6.4 Example

If the user intends on specifying the API key via an environment variable, then `dct-toolkit` provides a command to get the encrypted version of the API key:

```
./dct-toolkit encrypt_api_key apiKey=<DCT api key> dctUrl=<DCT instance URL>
```

13.7 Configure multiple DCT instances in `dct-toolkit`

13.7.1 Overview

This page describes a setup for a case where users need to connect to multiple DCT instances via `dct-toolkit` on a single engine. This would be useful in a case like upgrade testing. The user would upgrade the DCT instance in a lower environment (like a development environment) first, then run test APIs with `dct-toolkit`. Once all of the tests are successful, the user can upgrade the higher environment (like a production environment) and run checks via `dct-toolkit`. Being able to switch between DCT instances without hassle can be easily configured in `dct-toolkit`, as described below.

13.7.2 Example

The user will first need to create separate `dct-toolkit` configuration files for the environments that need to connect via `dct-toolkit`. This can be done by using the `configFileOutputLocation` option available for the `create_config` command:

```
./dct-toolkit create_config url=https://prod.... api_key=....  
configFileOutputLocation=<prod config file path>
```

```
./dct-toolkit create_config url=https://dev.... api_key=....  
configFileOutputLocation=<dev config file path>
```

Once the configuration files are created, users can create separate aliases for the config files generated above:

```
alias dct-toolkit-prod="DCT_TOOLKIT_CONFIG_FILE=<prod config file path>"  
alias dct-toolkit-dev="DCT_TOOLKIT_CONFIG_FILE=<dev config file path>"
```

Once the aliases are set up, users can then switch between production and development environments by invoking the corresponding aliases.

To get sources from the production environment, use the following command:

```
dct-toolkit-prod  
./dct-toolkit get_sources
```

To get sources from the development environment, use the following command:

```
dct-toolkit-dev  
./dct-toolkit get_sources
```

13.8 Logging

dct-toolkit has inbuilt logging support, which will export all of the logs to a default location under the user's home directory. The full path will be `<user home directory>/dct-toolkit/logs`. dct-toolkit creates a separate log file for every date.

Log level and log directory are both configurable and can be specified in the configuration file or via the environment variable. Please refer to the Installation and setup page for more information.

14 Developer resources

- [API requests and reporting](#) (see page 294)
- [API references](#) (see page 295)

14.1 API requests and reporting

14.1.1 Introduction

This article showcases example requests to the various data APIs supported by DCT.

DCT provides interactive API documentation that allows users to experiment with the APIs in their web browser. The interactive API documentation can be accessed by entering the hostname for DCT and the **/api** path into a browser's address bar. For example, if DCT is running on host `gateway.example.com`, then enter `https://gateway.example.com/api` into the browser's address bar.

To simplify development, Python and Go programming libraries are available. The **Python** bindings can be found on PyPi [here](#)⁹³. The latest version can be installed with the following command:

```
pip install delphix-dct
```

The **Go** bindings can be found on `go.dev` [here](#)⁹⁴.

14.1.2 Engines

This section showcases some examples of querying the Engines endpoint for information about connected Delphix Virtualization Engines. These examples leverage the generated Python bindings:

```
import delphix.api.gateway
import delphix.api.gateway.configuration
import delphix.api.gateway.api.management_api
cfg = delphix.api.gateway.configuration.Configuration()
cfg.host = "https://localhost/v2"

# For example purposes

cfg.verify_ssl = False

# Replace the string with your own API key

cfg.api_key['ApiKeyAuth'] = 'apk 3.tEd4DXFce'
```

⁹³ <https://pypi.org/project/delphix-dct-api/>

⁹⁴ <https://pkg.go.dev/github.com/delphix/dct-sdk-go>

```
api_client = delphix.api.gateway.ApiClient(configuration=cfg)
engines_api = delphix.api.gateway.api.management_api.ManagementApi(api_client)
print(engines_api.get_registered_engines())
```

The result should appear similar to the following:

```
{'items': [{ 'connection_status': 'ONLINE',
             'cpu_core_count': 2,
             'data_storage_capacity': 23404216320,
             'data_storage_used': 11589626880,
             'hostname': 'avm.delphix.com',
             'id': 1,
             'insecure_ssl': True,
             'memory_size': 8589934592,
             'name': 'vmname',
             'password': '*****',
             'status': 'CREATED',
             'tags': [],
             'type': 'UNSET',
             'unsafe_ssl_hostname_check': False,
             'username': 'admin',
             'uuid': 'ec2fbfea-928b-07f8-94c4-29fea614624f',
             'version': '6.1.0.0' } ] }
```

14.2 API references

The below API references are for DCT version 13.0.0. For the API references of previous versions, please visit the **API references** page in the corresponding version's documentation.

(Open API documentation is only available to view online)