

# **Hyperscale Compliance Home**

Hyperscale Compliance

Exported on 02/05/2024

## Table of Contents

1	Welcome to the Hyperscale Compliance documentation!.....	9
2	Quick references .....	15
3	Release notes.....	16
3.1	New features.....	16
3.1.1	17.0.0 release.....	16
3.1.2	16.0.0 release.....	16
3.1.3	15.0.0 release.....	16
3.1.4	14.0.0 release.....	16
3.1.5	13.0.0 release.....	17
3.1.6	12.0.0 release.....	17
3.1.7	11.0.0 release.....	18
3.1.8	10.0.0 release.....	18
3.1.9	9.0.0 release.....	18
3.1.10	8.0.0 release.....	18
3.1.11	7.0.0.0 release.....	19
3.1.12	6.0.0.0 release.....	19
3.1.13	5.0.0.1 releases.....	20
3.1.14	5.0.0.0 release.....	20
3.1.15	4.1.0 release.....	20
3.1.16	4.0.0 release.....	21
3.1.17	3.0.0.1 release.....	21
3.1.18	3.0.0.0 release.....	21
3.1.19	2.0.0.1 release.....	22
3.1.20	2.0.0 release.....	22
3.2	Fixed issues.....	22
3.2.1	Release 17.0.0.....	22
3.2.2	Release 16.0.0.....	23

3.2.3	Release 15.0.0.....	23
3.2.4	Release 14.0.0.....	23
3.2.5	Release 13.0.0.....	24
3.2.6	Release 12.0.0.....	24
3.2.7	Release 11.0.0.....	25
3.2.8	Release 10.0.0.....	25
3.2.9	Release 9.0.0.....	25
3.2.10	Release 8.0.0.....	25
3.2.11	Release 7.0.0.0.....	26
3.2.12	Release 6.0.0.0.....	26
3.2.13	Release 5.0.0.1.....	27
3.2.14	Release 5.0.0.0.....	27
3.2.15	Release 4.1.0.....	27
3.2.16	Release 4.0.0.....	28
3.2.17	Release 3.0.0.1.....	29
3.2.18	Release 3.0.0.....	29
3.3	Known issues.....	29
3.3.1	Release 17.0.0.....	30
3.3.2	Release 16.0.0.....	31
3.3.3	Release 15.0.0.....	32
3.3.4	Release 14.0.0.....	34
3.3.5	Release 13.0.0.....	35
3.3.6	Release 12.0.0.....	37
3.3.7	Release 11.0.0.....	39
3.3.8	Release 10.0.0.....	40
3.3.9	Release 9.0.0.....	41
3.3.10	Release 8.0.0.....	42
3.3.11	Release 7.0.0.0.....	43
3.3.12	Release 6.0.0.0.....	44

3.3.13	Release 5.0.0.0.....	45
3.3.14	Release 4.1.0.....	47
3.3.15	Release 4.0.0.....	49
3.3.16	Release 3.0.0.1.....	50
3.3.17	Release 3.0.0.....	52
3.3.18	Release 2.0.0.....	53
4	Overview.....	55
4.1	Hyperscale Compliance deployment architecture.....	55
4.1.1	Hyperscale Compliance components.....	55
4.1.1.1	Hyperscale Compliance Orchestrator.....	56
4.1.1.2	Staging area.....	56
4.1.1.3	Continuous Compliance Engine cluster.....	56
4.1.1.4	Source and target data sources.....	57
4.2	The Continuous Compliance platform.....	57
4.3	Next steps.....	58
5	Getting started.....	59
5.1	Hyperscale Compliance architecture.....	59
5.1.1	Controller service.....	59
5.1.2	Unload service.....	60
5.1.3	Masking service.....	60
5.1.4	Load service.....	61
5.2	Data source support.....	61
5.2.1	Oracle connector.....	61
5.2.1.1	Supported Data Types.....	62
5.2.1.2	Property values.....	62
5.2.1.3	Known limitations.....	63
5.2.2	MS SQL Connector.....	63
5.2.2.1	Supported versions.....	63
5.2.2.2	Supported data types.....	63

5.2.2.3	Property Values.....	64
5.2.2.4	Known Limitations .....	64
5.2.3	Delimited files connector .....	64
5.2.3.1	Prerequisites .....	64
5.2.3.2	Property values .....	65
5.2.3.3	Supported data types.....	66
5.2.3.4	Known limitations .....	66
5.2.4	MongoDB connector.....	67
5.2.4.1	Supported versions.....	67
5.2.4.2	Prerequisites .....	67
5.2.4.3	Property values .....	68
5.2.4.4	Known limitation: .....	69
5.2.5	Parquet connector .....	69
5.2.5.1	Prerequisites .....	69
5.2.5.2	Property values .....	71
5.2.5.3	Supported data types.....	71
5.2.5.4	Known limitations .....	72
5.3	Supported platforms.....	72
5.3.1	Supported Continuous Compliance/Data versions .....	72
5.3.2	Supported browsers (only API client) .....	73
5.4	Network requirements .....	73
5.5	Deployment .....	74
5.5.1	Docker compose .....	74
5.5.1.1	Host requirements (Docker compose) .....	75
5.5.1.2	Installation and setup (Docker compose) .....	76
5.5.1.3	Custom configuration.....	94
5.5.1.4	Upgrading the Hyperscale Compliance orchestrator (Docker compose) ...	95
5.5.1.5	How to generate a support bundle (Docker compose) .....	98
5.5.1.6	Managing the storage space .....	100

5.5.1.7	Migrating to Kubernetes .....	101
5.5.2	Kubernetes .....	106
5.5.2.1	Host requirements (Kubernetes) .....	106
5.5.2.2	Docker image registry and names for Hyperscale services.....	107
5.5.2.3	Installation and setup (Kubernetes) .....	109
5.5.2.4	Bootstrapping API keys (Kubernetes) .....	115
5.5.2.5	Hyperscale logs (Kubernetes).....	116
5.5.2.6	Limitations (Kubernetes).....	116
5.5.2.7	Upgrading the Hyperscale Compliance orchestrator (Kubernetes) .....	116
5.5.2.8	How to generate a support bundle (Kubernetes).....	118
5.5.3	Podman compose.....	119
5.5.3.1	Host requirements (Podman Compose) .....	120
5.5.3.2	Installation and setup (Podman Compose) .....	121
5.5.3.3	How to generate a support bundle (Podman Compose) .....	136
5.6	NFS server installation.....	138
5.6.1	Delphix Continuous Data Engine empty VDB .....	138
5.6.1.1	Continuous Data Engine installation.....	139
5.6.1.2	Discover and configure Hyperscale Compliance Orchestrator environment .....	139
5.6.1.3	Provision an empty VDB .....	140
5.6.2	NFS file server .....	141
5.6.3	Storage requirements for the NFS file server.....	141
5.7	Accessing the Hyperscale Compliance API .....	141
5.7.1	Authentication .....	142
5.7.2	Creating an API key.....	142
5.7.3	Using the newly generated key .....	143
5.7.4	Default API Version.....	143
6	How to setup a Hyperscale Compliance job .....	144
6.1	Pre-checks.....	144

6.2	API Flow to Setup a Hyperscale Compliance Job .....	144
6.3	Engines API .....	145
6.4	MountFileSystems API .....	146
6.5	ConnectorInfo API.....	146
6.6	StructuredDataFormat APIs .....	152
6.7	DataSets API .....	153
6.8	Jobs API .....	167
6.9	JobExecution API.....	171
7	How to Sync a Hyperscale Job .....	176
7.1	How to import a Job from Continuous Compliance Engine .....	176
7.2	How to re-import a Job from Continuous Compliance Engine .....	178
7.3	Script to automatically import/re-import a Job from Continuous Compliance Engine .....	178
7.4	How to sync global settings from a Delphix Continuous Compliance Engine .....	179
7.5	Limitations.....	180
8	How to cancel a Hyperscale job .....	181
9	Configuration settings .....	182
9.1	Commonly used properties .....	182
9.2	Other properties .....	186
10	Hyperscale Compliance API.....	200
10.1	Accessing the Hyperscale Compliance API.....	200
10.2	View the API reference .....	200
11	Cleaning up execution data.....	202
12	Hyperscale profilers.....	203
12.1	Parquet profiler .....	203
12.1.1	Installation and setup (Parquet profiler) .....	204
12.1.1.1	Pre-requisites .....	204
12.1.1.2	Procedure .....	204
12.1.2	Executing a profiler task.....	206

12.1.2.1 Perform the following steps to execute a profiler task .....	206
12.1.2.2 The DL entities within the default Profiler-Set with their algorithms .....	216





# 1 Welcome to the Hyperscale Compliance documentation!

When databases contain billions of rows of data, it can take weeks to protect sensitive data and PII using manual processes or bulk masking to anonymize the data. Hyperscale Compliance from Delphix provides incredibly fast masking speeds for large datasets enabling continuous compliant data delivery for CI/CD and DevOps initiatives.

Hyperscale Compliance does this by distributing the masking workload for a single job across multiple virtual Continuous Compliance Engines, reducing the time to mask large databases through increased scalability and efficiency.

This information explains how to deploy Hyperscale Compliance, 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 Hyperscale Compliance documentation versions in PDF format.

-  16.0.0\_HyperscaleCompliance.pdf [\(see page 9\)](#)
-  15.0.0\_HyperscaleCompliance.pdf [\(see page 9\)](#)

- [14.0.0\\_HyperscaleComplia...](#) (see page 9)
- [13.0.0\\_HyperscaleComplia...](#) (see page 9)
- [12.0.0\\_HyperscaleComplia...](#) (see page 9)

- [11.0.0\\_HyperscaleComplia...](#) (see page 9)
- [10.0.0\\_HyperscaleComplia...](#) (see page 9)
- [9.0.0\\_HyperscaleComplian...](#) (see page 9)

- [8.0.0\\_HyperscaleComplian...](#) (see page 9)
- [7.0.0\\_HyperscaleComplian...](#) (see page 9)
- [6.0.0\\_HyperscaleComplian...](#) (see page 9)

- [5.0.0\\_HyperscaleComplian...](#) (see page 9)
- [4.1.0\\_HyperscaleComplian...](#) (see page 9)
- [4.0.0\\_HyperscaleComplian...](#) (see page 9)

- [3.0.0\\_HyperscaleComplian...](#) (see page 9)
- [1.0.0\\_HyperscaleComplian...](#) (see page 9)

## 2 Quick references

- [Overview](#)<sup>1</sup>
- [Architecture](#)<sup>2</sup>
- [Data source support](#)<sup>3</sup>
- [New features](#)<sup>4</sup>
- [Fixed issues](#)<sup>5</sup>

---

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

2 <https://hyperscalemasking.delphix.com/docs/latest/hyperscale-compliance-architecture>

3 <https://hyperscalemasking.delphix.com/docs/latest/data-source-support>

4 <https://hyperscalemasking.delphix.com/docs/latest/new-features>

5 <https://hyperscalemasking.delphix.com/docs/latest/fixes>

## 3 Release notes

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

### 3.1 New features

#### 3.1.1 17.0.0 release

This release supports the following feature/features:

- **Addition of Apache Spark as data writer in Delimited Files connector**

The Delimited files connector now comes with PySpark (Apache Spark) to split the large files with added advantage. You can now select the backend data writer that works for your use case. To know more about these choices, refer to [Delimited Files Connector](#) (see page 61).

#### 3.1.2 16.0.0 release

This release supports the following feature/features:

- **Introduction of the Parquet connector**

This release introduces a Parquet connector that can be used to mask large Parquet files available on AWS S3 buckets. The connector splits the large files into smaller chunks, passes them to the masking service, and joins them back on the target location.

#### 3.1.3 15.0.0 release

This release supports the following feature/features:

- This release adds an option to configure whether you want to load empty values of Oracle BLOB/CLOB column as Null or empty string. A new property `load_empty_lobs_as_null` has been added under the 'target\_config' job to configure the same. It can be configured at the job level. The default value of this property is false (i.e., load as empty string).

#### 3.1.4 14.0.0 release

This release supports the following feature/features:

- **Oracle: Lookup references of a table recursively**

In this release, we have improved the pre-load process. Now, the pre-load process has been enhanced to perform recursive searches for table references until the final reference is found.



- **Support for masking structured data(XML/JSON) embedded in a database column**

This release introduces the support for masking the field values in XML / JSON data stored as CLOB in a database column. This has been achieved with the addition of a new entity called `structured-data-format` to the Hyperscale.

### 3.1.5 13.0.0 release

This release supports the following feature/features:

- **Job Cancellation - Phase 5**

This release provides an end-to-end Job Cancellation feature, offering you to cancel any Hyperscale Job Execution. All the running processes of Unload, Masking, Load, and Post Load Task will be canceled.

- **Introduction of MongoDB Connector**

This connector enhances data security by seamlessly masking extensive MongoDB collections. This ensures the continued usability of data while providing robust protection for sensitive information. The connector now offers enhanced export capabilities, allowing you to split collections based on their preferences. Data masking is seamlessly applied through the dedicated masking service, and the masked data is seamlessly imported into the designated target collection.

- This release introduces a separate artifact known as the MongoDB Profiler that can be downloaded from the same location as Hyperscale. It is an optional tool designed for profiling MongoDB collection data, generating an inventory of sensitive columns, and submitting the payload to the dataset API. The Profiler artifact includes a README file that provides detailed usage instructions. For information on the format of the dataset API payload, refer to the DataSets API section in this document.

- **New execution endpoint**

This release adds a new API GET endpoint ( `/executions/summary` ) to the existing JobExecution API to get the summarised overview of all the Executions of a particular Hyperscale Job.

### 3.1.6 12.0.0 release

This release supports the following feature/features:

- **Job cancellation - Phase 4**

This release adds the capability for users to cancel a Hyperscale Job Execution once the unload task of the execution has been finished. This will result in the cancellation of all ongoing processes related to masking and load tasks.

- **Introduction of Delimited Files Connector**

This release introduces a delimited files connector that can be used to mask large delimited flat files (with a delimiter of single character length) available on an NFS location. The connector splits the large into user-provided chunks, passes it to the masking service, and joins back.

- In this release, we're excluding pre-load and post-load processes for empty tables, leading to enhanced performance in scenarios where datasets contain empty tables.

### 3.1.7 11.0.0 release

This release supports the following feature/features:

- **Job cancellation - Phase 3**

With this release, you can cancel the MSSQL Hyperscale Job Execution while the load task of the execution is running.

- **Improvements in the Hyperscale job sync feature**

- This release introduces a new API endpoint `PUT /import/{datasetId}` to update the existing dataset and connector on the Hyperscale Compliance Orchestrator with refreshed ruleset from the Continuous Compliance Engine.
- This release provides a utility script to automate the process of creating/updating a dataset and connector at Hyperscale, by exporting a masking job from the Continuous Compliance engine.

For more details, refer to [How to Sync a Hyperscale Job](#) (see page 176) documentation.

- **Oracle NLS Support for Japanese character set**

This release adds Oracle NLS Support for the Japanese character set.

### 3.1.8 10.0.0 release

This release supports the following feature/features:

- **Job cancellation - Phase 2**

With this release, you can cancel an Oracle Hyperscale Job Execution while the Load task of the execution is running. This feature is not available for MSSQL connectors.

### 3.1.9 9.0.0 release

This release supports the following feature/features:

- **Job cancellation - Phase 1**

With this release, you can cancel a Hyperscale Job Execution while the Post Load task of the execution is running. For more details, refer to the [How to Cancel a Hyperscale Job](#)<sup>6</sup> documentation.

### 3.1.10 8.0.0 release

This release supports the following feature/features:

- **Support for Kubernetes deployment**

This release introduces the capability to deploy the Hyperscale Compliance Orchestrator on a single-

---

<sup>6</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5309804/%289.0.0%29+How+to+cancel+a+Hyperscale+job>

node Kubernetes cluster by using the helm charts. For more details, refer to the [\(8.0.0\) Installation and setup \(Kubernetes\)](#)<sup>7</sup> documentation.

- **Enhanced post-load task status**

This release provides you with comprehensive visibility into the progress status of the post-load task (for example, ongoing process for constraints, indexes, and triggers) using the execution API endpoints. For more details, refer to the [How to Setup a Hyperscale Compliance Job](#) documentation.

- **Oracle connector - degree of parallelism for recreating indexes**

This release provides you the ability to specify and configure the degree of parallelism(DOP) per Oracle job to recreate the index in the post-load process. Currently, the recreate index DDL utilizes only the default Degree of Parallelism set by the oracle but now you can specify the custom value that can enhance the performance of the index recreation process. For more details, refer to the [\(8.0.0\) Hyperscale Compliance API](#)<sup>8</sup> documentation.

- **Introduction of semantic versioning (Major.Minor.Patch)**

With this release, Hyperscale introduced support for Kubernetes deployment through helm charts. As helm charts support 3 part Semantic Versioning, hence release 8.0.0 onwards Hyperscale will also follow the 3 part Semantic Versioning instead of 4 parts semantic versioning.

### 3.1.11 7.0.0.0 release

This release supports the following feature/features:

- **Performance improvement**

This release introduces impactful changes aimed at enhancing the performance of the masking task within the Hyperscale Job, ultimately resulting in improved overall job performance. The following key changes have been implemented:

- Changes in Masking Service to increase the Compliance Engine utilization by Hyperscale.
- Masking Service will no more process tables having 0 rows.

- **Oracle**

This release supports tables with subpartition indexes during load.

### 3.1.12 6.0.0.0 release

This release supports the following feature/features:

- **New file based endpoints**

- file-download: This release introduces a new API endpoint to download the execution and dataset responses. For more information, refer to the [\(6.0.0\) Hyperscale Compliance API](#)<sup>9</sup> documentation.
- file-upload: This release introduces a new API endpoint to upload a file that currently can be used to create or update the dataset using POST /data-sets/file-upload and PUT /data-sets/

<sup>7</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5342943>

<sup>8</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5408602/%288.0.0%29+Hyperscale+Compliance+API>

<sup>9</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5343436/%286.0.0%29+Hyperscale+Compliance+API>

file-upload/{dataSetId} endpoints. For more information, refer to the [\(6.0.0\) Hyperscale Compliance API](#)<sup>10</sup> documentation.

- **MSSQL database scalability improvements**

This release improves the overall job performance by adding the handling of primary key constraints.

### 3.1.13 5.0.0.1 releases

5.0.0.1 is a patch release specifically aimed at addressing a critical bug. For more information, see [\(5.0.0\) Fixed issues](#)<sup>11</sup>.

### 3.1.14 5.0.0.0 release

This release supports the following feature/features:

- **MS SQL connector**

This release adds the MS SQL connector implemented as separate services that include unload and load services. These connector services enable Hyperscale Compliance for MS SQL databases.

- **New execution endpoints**

This release adds a new API GET endpoint ( `/executions/{id}/summary` ) to the existing JobExecution API to get the overview of the progress of a Job Execution. In addition to this, the existing `GET /executions/{id}` endpoint has been extended to have additional filters based on the task name and the task's metadata object's status. For more information, refer to the Execution API in the [\(5.0.0\) Hyperscale Compliance API](#)<sup>12</sup> section.

- **Multi-column algorithm support**

With this release, Multi-Column Algorithms can also be provided in the `/data-sets` endpoint. For more information, refer to the Dataset API in the [\(5.0.0\) Hyperscale Compliance API](#)<sup>13</sup> section. Additionally, existing Continuous Compliance jobs containing multi-column algorithm-related fields can now be imported `via/import` endpoint.

### 3.1.15 4.1.0 release

This release supports the following feature/features:

- **Capability to limit the number of connections**

This release adds the capability to limit the number of connections to the source and target databases using the new API parameters as `Max_concurrent_source_connection` and `Max_concurrent_target_connection` under new `source_configs` and `target_configs` respectively. Using this property, you can fine-tune the number of connections as per source target infra to get better performance. For more information, refer to the [\(4.1.0\) Hyperscale Compliance API](#)<sup>14</sup> documentation.

---

<sup>10</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5343436/%286.0.0%29+Hyperscale+Compliance+API>

<sup>11</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5310675/%285.0.0%29+Fixed+issues>

<sup>12</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5410053/%285.0.0%29+Hyperscale+Compliance+API>

<sup>13</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5410053/%285.0.0%29+Hyperscale+Compliance+API>

<sup>14</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5245694/%284.1.0%29+Hyperscale+Compliance+API>

- **Increased API object limit**

This release increases the API object limit from 1000 to 10000.

### 3.1.16 4.0.0 release

This release supports the following feature/features:

- **Hyperscale job sync**

This release adds the ability to:

- Import masking jobs inventory from Continuous Compliance engines into connector and dataset info of Hyperscale Compliance Orchestrator with the sync [\(4.0.0\) Accessing the Hyperscale Compliance API](#)<sup>15</sup> endpoint.
- Import global settings that include Algorithms/Domains from Continuous Compliance Engines to Hyperscale Clustered Continuous Compliance Engines using the sync [\(4.0.0\) Accessing the Hyperscale Compliance API](#)<sup>16</sup> endpoint.

For more information, refer to the [\(4.0.0\) How to sync a Hyperscale job](#)<sup>17</sup> section.

- **Add configuration properties through .env file**

This release adds an additional capability to override commonly used configuration properties through the .env file. You can now update application properties in this file before starting the application. For more information, refer to the [\(4.0.0\) Configuration settings](#)<sup>18</sup> section.

### 3.1.17 3.0.0.1 release

3.0.0.1 is a patch release specifically aimed at addressing critical bugs. For more information, see [\(3.0.0\) Fixed issues](#)<sup>19</sup>.

### 3.1.18 3.0.0.0 release

This release supports the following feature/features:

- **Oracle connector**

This release includes the Oracle connector implemented as separate services, including unload and load services. These connector services enable Hyperscale Compliance for Oracle databases.

- **Parallel processing of tables**

This release processes all tables provided through the data-set API in parallel through the four operational stages - unload, masking, upload, and post-load to minimize the total time it takes to mask the complete data set.

---

<sup>15</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5410780/%284.0.0%29+Accessing+the+Hyperscale+Compliance+API>

<sup>16</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5410780/%284.0.0%29+Accessing+the+Hyperscale+Compliance+API>

<sup>17</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5311479/%284.0.0%29+How+to+sync+a+Hyperscale+job>

<sup>18</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5311507/%284.0.0%29+Configuration+settings>

<sup>19</sup> <https://delphixdocs.atlassian.net/wiki/spaces/HC/pages/5410967/%283.0.0%29+Fixed+issues>

- **Monitoring**  
This release provides monitoring APIs so that you can track the progress of tables in your data set through the unload, masking, upload, and post-load phases. This API also provides a count of rows being processed through different stages.
- **Restartability**  
This release includes the ability to restart a failed process.
- **Clean up**  
This release supports cleaning data from previous job execution.

### 3.1.19 2.0.0.1 release

2.0.0.1 is a patch release specifically aimed at addressing critical bugs and has the following updates:

- Upgraded spring boot version to 2.5.12.
- Minor view-only changes in swagger-based API client.

### 3.1.20 2.0.0 release

2.0.0 is the initial release of Hyperscale Compliance. Hyperscale Compliance is an API-based interface that is designed to enhance the performance of masking large datasets. It allows you to achieve faster masking results using the existing Delphix Continuous Compliance offering without adding the complexity of configuring multiple jobs.

## 3.2 Fixed issues

This section describes the issues fixed in Hyperscale Compliance.

### 3.2.1 Release 17.0.0

Key	Summary
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error
HM-2962	Incorrect gathering target table group list in case of the table referred from other tables.
HSC-176	Delimited connector: values with data type double and int64 will always be converted into a string
HSC-177	Delimited connector: Irrespective of user provided enclosure character, the output strings will be quoted using double quotes

Key	Summary
HSC-279	Wrong counts in Unload: Using HS 14 delimited file connector, while masking two files in the dataset, job shows failed status, because the counts are all wrong
HSC-286	Load failed because the SQLite3 database locked up. Masking 15 files with the same data set. Unload has the wrong count.
HSC-361	Hyperscale delimited connector should not add double quotes to output file

### 3.2.2 Release 16.0.0

Key	Summary
HM-2748	Unable to upload sync bundle of more than 20 MB

### 3.2.3 Release 15.0.0

Key	Summary
HM-2505	ServiceUnavailableException in Masking Service as numeric value out of range
HM-2554	Missing security context causes controller to fail with permission denied for JDBC connection to SQLite DB
HM-2609	Configuration added to load empty values of oracle CLOB/BLOB columns as null or empty string

### 3.2.4 Release 14.0.0

Key	Summary
HSC-182	400 Bad Request from POST <a href="http://unload-service:8080/api/unload">http://unload-service:8080/api/unload</a>
HM-2413	NTRS: Pre Load failed with errors: Error executing statement: Error disabling constraint

Key	Summary
HM-2505	ServiceUnavailableException in Masking Service as numeric value out of range
HM-2554	Missing security context causes controller to fail with permission denied for JDBC connection to SQLite DB
HM-2568	Hyperscale API's working without auth key in k8s setup

### 3.2.5 Release 13.0.0

Key	Summary
HSC-178	Delimited load service will show the status as FAILED while it is waiting for all split files to be masked in order to perform join.
HSC-179	While executing masking against a 1.1 TB second time, the load job failed mid-execution.
HM-2399	MSSQL: The lower bound of partitioning column is larger than the upper bound.
HM-2443	Masking service is mapping the wrong file metadata when multiple data_info objects with varying delimiters are passed during data-sets creation.

### 3.2.6 Release 12.0.0

Key	Summary
HM-2268	Add support to configure FileMetadata related characters for structured (i.e. XML) data
HM-2344	sync-compliance-engine endpoint failing with "Invalid input" error
HM-2399	MSSQL: The lower bound of partitioning column is larger than the upper bound.



### 3.2.7 Release 11.0.0

There are no fixed issues in this release.

### 3.2.8 Release 10.0.0

Key	Summary
HM-1360	Job is created with default value of 'retain_execution_data' parameter if an invalid value for the same parameter is passed while creating the job object
HM-2041	Capture high-level processing timings for each process for each object
HM-2084	cancel script (cancel.sh ) does not read values from .env file
HM-2199	Masking service paginated GET call, skip the result of last page

### 3.2.9 Release 9.0.0

Key	Summary
HM-1845	java.sql.SQLException: Protocol violation While fetching metadata of the table
HM-1980	Post load status incorrectly reporting consolidated start and end timing.
HM-2017	The masking process is stuck in running when a newly registered engine is used in execution after the timeout time of the Engine

### 3.2.10 Release 8.0.0

Key	Summary
HM-1606	When tables have more than 100 columns on Oracle, the default API page at 100 causes an issue where we don't match the column name via the masking API

Key	Summary
HM-1787	Oracle: Load service holds a lock on class level, causing parallel jobs stuck in the flow of preload
HM-1814	MSSQL: Load service holds the lock on class level, causing parallel jobs stuck in the flow of preload

### 3.2.11 Release 7.0.0.0

Key	Summary
HM-1617	Oracle - Hyperscale engine fails to mark partitions of the partitioned index as unusable in pre-load and fails to rebuild a partitioned index in post load leaving indexes in an unusable state
HM-1684	Running multiple Hyperscale jobs in parallel, using the same set of Masking Engines and the same value of env_name_prefix and app_name_prefix can cause job failure with the missing ruleset, any other masking objects, or execution error.

### 3.2.12 Release 6.0.0.0

Key	Summary
HM-1513	MSSQL - Slowness while performing load with large tables( more than 4M rows and 10 Columns)
HM-1521	MSSQL - Post load fails with 'transaction log is full due to 'ACTIVE_TRANSACTION' for large tables
HM-1608	SQLCLR doesn't load the data when the table has enabled triggers owned by a different user
HM-1645	Failed clean up of previous execution causes "\"File Format already exists with identifier: HYPERSCALE_16_9fd11fae45861fdb18fb658fb950ceab.fmt\" issue for next execution of same job
HM-1656	nginx configuration client max body size limits the size of the payload to post for the dataset

### 3.2.13 Release 5.0.0.1

Key	Summary
HM-1 561	Oracle Load Failure: SQL loader control files doesn't contain character length when column size is less than 256 CHAR

### 3.2.14 Release 5.0.0.0

Key	Summary
HM-9 71	For HTTP protocol-type requests, trust store fields are accepted and displayed in response.
HM-1 472	Oracle Unload service doesn't release the lock if fails to initialize the connection pool and the job stuck in a running state
HM-1 520	Masking service: Starting execution throwing NPE after container restart
HM-1 522	Update Oracle JDBC Driver to 21.3.0.0

### 3.2.15 Release 4.1.0

Key	Summary
HM-1155	Diagnosability: How do I tell which masking job on the masking engine relates to the failed message on the HS Jobs API status
HM-1168	The error text in Post Load failures is misleading/unclear
HM-1191	Optimization: We should execute Select Count (*) in parallel to the Oracle Unload process. It could take a significant amount of time to count data in large tables as well as 1000 objects.

Key	Summary
HM-1201	Unable to import a ruleset with 5000 tables to HS 4.0, getting an error message.
HM-1210	While trying to process an HS job for 5000 table schema, ran out of Oracle cursor, and then the SQLite database locked up.
HM-1265	Oracle - Any index on a column having no constraint on it, is not getting dropped
HM-1334	Can't drop the index, because the index is not owned by the user we used to connect.
HM-1378	Oracle - Load service needs to include index owner name when attempting to modify/rebuild partition indexes owned by different db user

### 3.2.16 Release 4.0.0

Key	Summary
HM-177	Able to POST /hyperscale-masking/jobs with min job memory > max job memory
HM-530	POST/PUT request dataSet API error response received with empty/missing/invalid 'source'/'target' object/values can be improved
HM-754	POST/PUT connector jdbc_url, username, and password should be mandatory for Oracle load and unload service
HM-789	Error message upon not setting the 'SSL' field to False indicates 'insecure_ssl' property which no longer exists in the schema
HM-858	Status of sub-task coming wrong when overall execution failed
HM-932	Suppress the password message for the controller log

Key	Summary
HM-1138	The description in the swagger doesn't match the API call
HM-1140	The error needs more information to diagnose a connector issue.

### 3.2.17 Release 3.0.0.1

Key	Summary
HM-858	Status of sub-task coming wrong when overall execution failed
HM-873	Intermittently there is a mismatch in loaded_rows displayed in the load task vs the actual rows loaded in the target table
HM-915	Load: driver support plugin throws ORA-02297: cannot disable constraint - dependencies exist error for foreign key

### 3.2.18 Release 3.0.0

Key	Summary
HM-294	The updated file format is not POST'ed on the Continuous Compliance Engine if the file format name is the same

## 3.3 Known issues

This section describes the known issues in Hyperscale Compliance.

### 3.3.1 Release 17.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None

Key	Summary	Workaround
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.
HM-2485	Cancel job will terminate more promptly in comparison with the earlier release. In certain situations, like unloading a very large number of MSSQL tables, there might not be a noticeable improvement.	None
HM-2593	MSSQL - In PreLoad table references not fetched till last level in relationship hierarchy	None
HM-2663	Cleanup docker-compose-sample.yaml file in the docker-compose tar bundle	None

### 3.3.2 Release 16.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None

Key	Summary	Workaround
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.
HM-2485	Cancel job will terminate more promptly in comparison with the earlier release. In certain situations, like unloading a very large number of MSSQL tables, there might not be a noticeable improvement.	None
HM-2593	MSSQL - In PreLoad table references not fetched till last level in relationship hierarchy	None
HM-2663	Cleanup docker-compose-sample.yaml file in the docker-compose tar bundle	None

### 3.3.3 Release 15.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None



Key	Summary	Workaround
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.
HM-2485	Cancel job will terminate more promptly in comparison with the earlier release. In certain situations, like unloading a very large number of MSSQL tables, there might not be a noticeable improvement.	None

Key	Summary	Workaround
HM-2593	MSSQL - In PreLoad table references not fetched till last level in relationship hierarchy	None
HM-2663	Cleanup docker-compose-sample.yaml file in the docker-compose tar bundle	None

### 3.3.4 Release 14.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None

Key	Summary	Workaround
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2413	In PreLoad table references not fetched till last level in relationship hierarchy	None
HM-2485	Cancel job will terminate more promptly in comparison with the earlier release. In certain situations, like unloading a very large number of MSSQL tables, there might not be a noticeable improvement.	None
HM-2663	Cleanup docker-compose-sample.yaml file in the docker-compose tar bundle	None

### 3.3.5 Release 13.0.0

#### Delimited Files Connector

Key	Summary	Workaround
HSC-174	Delimited unload service does not show incremental unload row count, it always should unloaded count is the same as the total count	None
HSC-176	Delimited connector: values with data type double and int64 will always be converted into a string	None
HSC-177	Delimited connector: Irrespective of user provided enclosure character, the output strings will be quoted using double quotes	None

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.

Key	Summary	Workaround
HM-2413	In PreLoad table references not fetched till last level in relationship hierarchy	None
HM-2485	Cancel job will terminate more promptly in comparison with the earlier release. In certain situations, like unloading a very large number of MSSQL tables, there might not be a noticeable improvement.	None

### 3.3.6 Release 12.0.0

#### Delimited Files Connector

Key	Summary	Workaround
HSC-174	Delimited unload service does not show incremental unload row count, it always should unloaded count is the same as the total count	None
HSC-176	Delimited connector: values with data type double and int64 will always be converted into a string	None
HSC-177	Delimited connector: Irrespective of user provided enclosure character, the output strings will be quoted using double quotes	None
HSC-178	Delimited load service will show the status as FAILED while it is waiting for all split files to be masked in order to perform join	None
HM-2443	Masking service is mapping the wrong file metadata when multiple data_info objects with varying delimiters are passed during data-sets creation.	Create individual data-sets for files with different delimiter character.

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.

### 3.3.7 Release 11.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None

Key	Summary	Workaround
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.

### 3.3.8 Release 10.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	None
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	Manually enable the unusable indexes
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None



Key	Summary	Workaround
HM-1929	Postload task status shows the old error message for the error field after restart the execution.	None
HM-2185	Post Load constraints/indexes/triggers queries keeps on running at target DB even after cancellation is completed at Hyperscale end.	Refresh the target database before restarting/rerunning the Job.

### 3.3.9 Release 9.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	Use a filter key which does not have any Null values
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle: Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None
HM-1397	Oracle Load fails for tables having triggers only with SQL*Loader-937 error	None

Key	Summary	Workaround
HM-1463	MSSQL - Load Service fails when the Table name contains '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restarting the execution.	None

### 3.3.10 Release 8.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	Use a filter key which does not have any Null values
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle: Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None

Key	Summary	Workaround
HM-1397	Oracle Load fails for tables having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when the Table name contains ' '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None
HM-1929	Postload task status shows the old error message for the error field after restarting the execution.	None

### 3.3.11 Release 7.0.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL - The job is stuck in a running state while using a filter key having NULL values.	Use a filter key which does not have any Null values
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None

Key	Summary	Workaround
HM-1382	Oracle: Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None
HM-1397	Oracle Load fails for tables having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when the Table name contains ' '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.12 Release 6.0.0.0

Key	Summary	Workaround
HM-663	Oracle: The load process is failing with "Error disabling constraint" for identity columns	None
HM-812	Application on the registered masking engine is not deleted with the cleanup	None
HM-1196	MSSQL Job is stuck in a running state while using a filter key having NULL values.	Use a filter key which does not have any Null values
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1382	Oracle: Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None

Key	Summary	Workaround
HM-1397	Oracle: Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None
HM-1463	MSSQL - Load Service fails when the Table name contains ' '	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.13 Release 5.0.0.0

Key	Summary	Workaround
HM-291	Hyperscale job execution with an intelligent load balancer configured is stuck in a loop if the job's max memory is more than totalAllocatedMemoryForJobs	Change the max memory to a value under the value of <code>totalAllocatedMemoryForJobs</code> the property configured on the Continuous Compliance Engine.
HM-652	Job execution is stuck in a running state if the mount server is powered off	Check the health of the mount server before starting a job.
HM-663	Oracle: Load process is failing with "Error disabling constraint" for identity columns	None
HM-718	Not all data on the mount server is cleaned up if the continuous compliance engine is stopped	Cleanup up the data manually from the mount server.
HM-745	The table name is not present in the error message while enabling/disabling triggers, indexes, constraints	Check the logs in container logs to get table details
HM-812	Application on the registered masking engine is not deleted with the cleanup	None

Key	Summary	Workaround
HM-817	Intermittently job fails with ORA-02270: no matching unique or primary key for this column-list	Restart the job using <code>PUT / executions/{id}/restart</code> and it will succeed.
HM-821	Hyperscale job does not handle post-load task properly during restart if failed in pre-load (disabling trigger/indexes/constraints) steps	After job execution is completed successfully, check and manually enable the disabled constraints.
HM-1196	MSSQL Job is stuck in a running state while using a filter key having NULL values.	Use a filter key which does not have any Null values
HM-1239	MSSQL- Unload fails for a schema or table name having a ']' character in it	None
HM0-1240	MSSQL - Unload fails for a column having '.' in its name	None
HM-1246	MSSQL - Unload fails with UnknownFormatConversionException for a table having % in its name	None
HM-1251	MSSQL: Row is not loaded if masked BIT type column has values other than true(1),false(0) or NULL	None
HM-1366	NPE displayed in hyperscale masking service logs just after masking task is done	None
Hm-1382	Oracle : Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None
HM-1397	Oracle Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1463	MSSQL - Load Service fails when Table name contains '	None
HM-1512	received_objects in Load step are more than succeeded_objects in Masking step intermittently	None

Key	Summary	Workaround
HM-1513	MSSQL - Slowness while performing load with large tables( more that 4M rows and 10 Columns)	None
HM-1521	MSSQL - Post load fails with 'transaction log is full due to 'ACTIVE_TRANSACTION' for large tables	None
HM-1523	MSSQL - Job fails while loading masked VARBINARY data	None
HM-1528	Initial delay in updating response of unload/masking/ load execution objects with large number of tables	None
HM-1561	Oracle Load Failure: sql loader control files doesn't contain character length when column size is less than 256 CHAR	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.14 Release 4.1.0

Key	Summary	Workaround
HM-291	Hyperscale job execution with intelligent load balancer configured is stuck in a loop if job's max memory is more than totalAllocatedMemoryForJobs	Change the max memory to a value under the value of <code>totalAllocatedMemoryForJobs</code> property configured on Continuous Compliance Engine.
HM-652	Job execution is stuck in running state if mount server is powered off	Check the health of mount server before starting a job.
HM-663	Load process is failing with "Error disabling constraint" for identity columns	None
HM-718	Not all data on mount server is cleaned up if masking engine is stopped	Cleanup up the data manually from the mount server.

Key	Summary	Workaround
HM-745	Table name is not present in error message while enabling/disabling triggers,indexes,constraints	Check the logs in container logs to get table details
HM-812	Application on registered masking engine is not deleted with cleanup	None
HM-817	Intermittently job fails with ORA-02270: no matching unique or primary key for this column-list	Restart the job using <code>PUT / executions/{id}/restart</code> and it will succeed.
HM-821	Hyperscale job does not handle post load task properly during restart if failed in pre-load (disabling trigger/indexes/constraints) steps	After job execution is completed successfully, check and manually enable the disabled constraints.
HM-1155	Diagnosibility: How do I tell which masking job on the masking engine relates to the failed message on the HS Jobs API status	Check the error details in masking service logs
HM-1168	The error text is inaccurate, and doesn't contain enough information to diagnose it without accessing logs on the Hyperscale server.	Check the error details in the logs
HM-1561	Oracle Load Failure: sql loader control files doesn't contain character length when column size is less than 256 CHAR	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None



### 3.3.15 Release 4.0.0

Key	Summary	Workaround
HM-291	Hyperscale job execution with intelligent load balancer configured is stuck in a loop if job's max memory is more than <code>totalAllocatedMemoryForJobs</code>	Change the max memory to a value under the value of <code>totalAllocatedMemoryForJobs</code> property configured on Continuous Compliance Engine.
HM-652	Job execution is stuck in running state if mount server is powered off	Check the health of mount server before starting a job.
HM-663	Load process is failing with "Error disabling constraint" for identity columns	None
HM-718	Not all data on mount server is cleaned up if masking engine is stopped	Cleanup up the data manually from the mount server.
HM-745	Table name is not present in error message while enabling/disabling triggers,indexes,constraints	Check the logs in container logs to get table details
HM-812	Application on registered masking engine is not deleted with cleanup	None
HM-817	Intermittently job fails with ORA-02270: no matching unique or primary key for this column-list	Restart the job using <code>PUT / executions/{id}/restart</code> and it will succeed.
HM-821	Hyperscale job does not handle post load task properly during restart if failed in pre-load (disabling trigger/indexes/constraints) steps	After job execution is completed successfully, check and manually enable the disabled constraints.
HM-1366	NPE displayed in hyperscale masking service logs just before cleanup is performed	None
HM-1382	Dataset having any one entry with invalid schema leaves indexes of other tables as UNUSABLE	None

Key	Summary	Workaround
HM-1397	Load fails for table having triggers only with SQL*Loader-937 error	None
HM-1561	Oracle Load Failure: sql loader control files doesn't contain character length when column size is less than 256 CHAR	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.16 Release 3.0.0.1

Key	Summary	Workaround
HM-177	Able to POST /hyperscale-masking/jobs with min job memory > max job memory	Change the max job memory value to higher than min job memory in API request.
HM-291	Hyperscale job execution with intelligent load balancer configured is stuck in a loop if job's max memory is more than totalAllocatedMemoryForJobs	Change the max memory to a value under the value of <code>totalAllocatedMemoryForJobs</code> property configured on Continuous Compliance Engine.
HM-652	Job execution is stuck in running state if mount server is powered off	Check the health of mount server before starting a job.
HM-663	Load process is failing with "Error disabling constraint" for identity columns	None
HM-684	Hypescale does not support other TIMESTAMP(6) datatype variations apart from TIMESTAMP	None
HM-718	Not all data on mount server is cleaned up if batch masking service is stopped	Cleanup up the data manually from mount server.

Key	Summary	Workaround
HM-745	Table name is not present in error message while enabling/disabling triggers,indexes,constraints	Check the logs in container logs to get table details.
HM-754	Able to POST/PUT a connector with whitespace as jdbc_url, username, password	Remove white space and use valid values for jdbc_url, username and password.
HM-789	Error message upon not setting 'ssl' field to False indicates 'insecure_ssl' property which no longer exists in the schema	None
HM-812	Application on registered masking engine is not deleted with cleanup	None
HM-817	Intermittently job fails with ORA-02270: no matching unique or primary key for this column-list	Restart the job using <code>PUT / executions/{id}/restart</code> and it will succeed.
HM-821	Hyperscale job does not handle post load task properly during restart if failed in pre-load (disabling trigger/indexes/constraints) steps	After job execution is completed successfully, check and manually enable the disabled constraints.
HM-935	Load service fails when source DB contains BLOB type data that is not simple text file data	None
HM-1561	Oracle Load Failure: sql loader control files doesn't contain character length when column size is less than 256 CHAR	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.17 Release 3.0.0

Key	Summary	Workaround
HM-177	Able to POST /hyperscale-masking/jobs with min job memory > max job memory	Change the max job memory value to higher than min job memory in API request.
HM-291	Hyperscale job execution with intelligent load balancer configured is stuck in a loop if job's max memory is more than totalAllocatedMemoryForJobs	Change the max memory to a value under the value of totalAllocatedMemoryForJobs property configured on Continuous Compliance Engine.
HM-652	Job execution is stuck in running state if mount server is powered off	Check the health of mount server before starting a job.
HM-663	Load process is failing with "Error disabling constraint" for identity columns	None
HM-684	Hypescale does not support other TIMESTAMP(6) datatype variations apart from TIMESTAMP	None
HM-718	Not all data on mount server is cleaned up if batch masking service is stopped	Cleanup up the data manually from mount server.
HM-745	Table name is not present in error message while enabling/disabling triggers,indexes,constraints	Check the logs in container logs to get table details.
HM-754	Able to POST/PUT a connector with whitespace as jdbc_url,username,password	Remove white space and use valid values for jdbc_url, username and password.
HM-789	Error message upon not setting 'ssl' field to False indicates 'insecure_ssl' property which no longer exists in the schema	None
HM-812	Application on registered masking engine is not deleted with cleanup	None

Key	Summary	Workaround
HM-817	Intermittently job fails with ORA-02270: no matching unique or primary key for this column-list	Restart the job using <code>PUT / executions/{id}/restart</code> and it will succeed.
HM-821	Hyperscale job does not handle post load task properly during restart if failed in pre-load (disabling trigger/indexes/constraints) steps	After job execution is completed successfully, check and manually enable the disabled constraints.
HM-858	Status of sub task coming wrong when overall execution failed	None
HM-873	Intermittently there is a mismatch in loaded_rows displayed in load task vs the actual rows loaded in target table	None
HM-915	Load: driver support plugin throws ORA-02297: cannot disable constraint - dependencies exist error for foreign key	None
HM-935	Load service fails when source DB contains BLOB type data that is not simple text file data	None
HM-1561	Oracle Load Failure: sql loader control files doesn't contain character length when column size is less than 256 CHAR	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

### 3.3.18 Release 2.0.0

The following is a list of the known issues in the Hyperscale Compliance version 2.0.0.

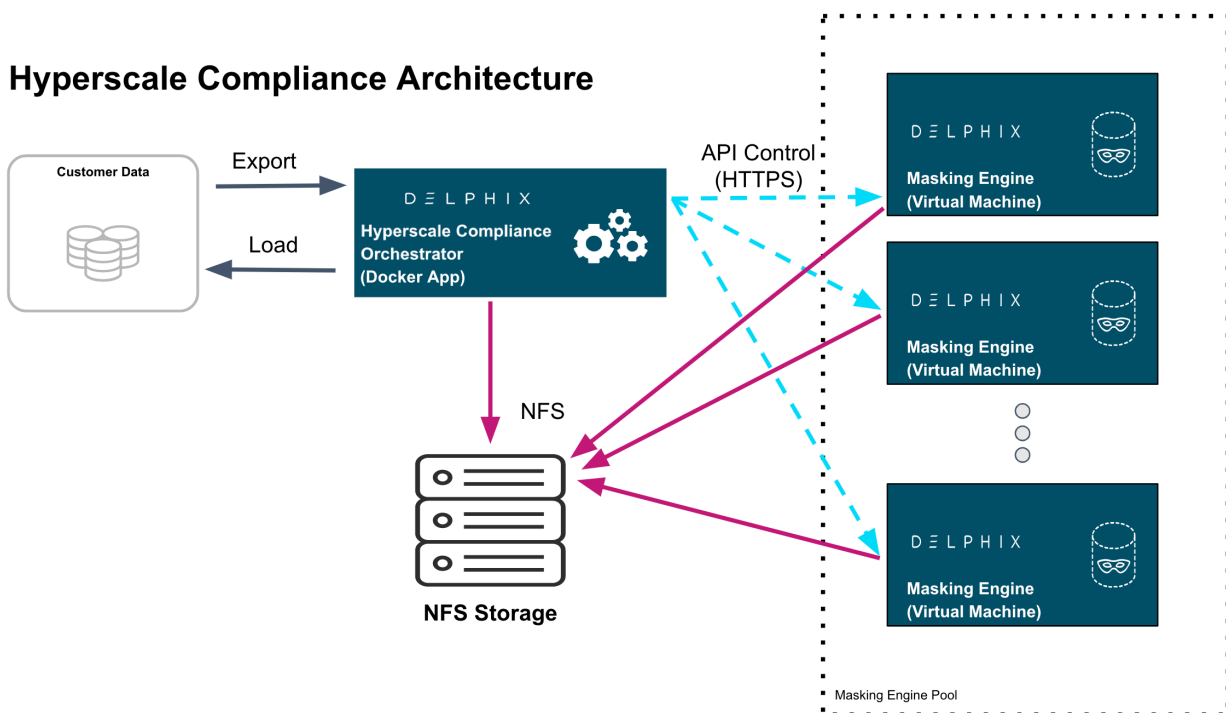
Key	Summary	Workaround
HM-294	Updated file format is not POST'ed on the Continuous Compliance Engine if file format name is same	<ul style="list-style-type: none"> <li>• After modifying the file format content, rename the file name of the file format.</li> <li>• Delete the existing uploaded file format from the attached Continuous Compliance Engines before executing the Hyperscale job with an updated file format.</li> </ul>
HM-291	Hyperscale job execution with intelligent load balancer configured is stuck in a loop if job's max memory is more than totalAllocatedMemoryForJobs	None
HM-216	POST/PUT /data-sets accepts duplicate values as source_files path in Single File Info Object	None
HM-177	Able to POST /hyperscale-masking/jobs with min job memory > max job memory	None
HM-1705	Improper error message in Hyperscale status response if CCE gets mount file system connection error	None

## 4 Overview

Hyperscale Compliance is an API-based interface that is designed to enhance the performance of masking large datasets. It allows you to achieve faster masking results using the existing Delphix Continuous Compliance offering without adding the complexity of configuring multiple jobs. Hyperscale Compliance first breaks the large and complex datasets into numerous modules and then orchestrates the masking jobs across multiple Continuous Compliance Engines. In general, datasets larger than 10 TB in size will see improved masking performance when run on the Hyperscale architecture.

### 4.1 Hyperscale Compliance deployment architecture

For achieving faster masking results, Hyperscale Compliance uses bulk import or export utilities of data sources. Using these utilities, it exports the data into smaller chunks of delimited files. The Hyperscale Compliance Orchestrator then configures the masking jobs of all the respective chunks across multiple Continuous Compliance Engines. Upon successful completion of the masking jobs, the masked data is imported back into the database.



#### 4.1.1 Hyperscale Compliance components

The Hyperscale Compliance architecture consists of four components mainly; the Hyperscale Compliance Orchestrator, Source/Target Connectors, the Continuous Compliance Engine Cluster, and the Staging Server.

### 4.1.1.1 Hyperscale Compliance Orchestrator

The Hyperscale Compliance Orchestrator is responsible for unloading the data from the source and horizontally scaling the masking process by initiating multiple parallel masking jobs across nodes in the Continuous Compliance Engine cluster. Once data is masked, it loads it back to the target data sources. Depending on the number of nodes in the cluster, you can increase or decrease the total throughput of an individual masking job. In the case of relational databases as source and target data sources, it also handles the pre-load (disabling indexes, triggers, and constraints) and post-load (enabling indexes, triggers, and constraints) tasks like disabling and enabling indexes, triggers, and constraints. Currently, the Hyperscale Compliance Orchestrator supports the following two strategies to distribute the masking jobs across nodes available :

- **Intelligent Load Balancing (Default):** This strategy considers each Continuous Compliance Engine's current capacity before assigning any masking jobs to the node Continuous Compliance Engines. It calculates the capacity using available resources on node Continuous Compliance Engines and already running masking jobs on the engines. Below is the formula used to calculate the capacity of the Continuous Compliance Engines:

```
Engine's current jobCapacity = Engine's total jobCapacity - no of currently running jobs on Engine
```

```
Engine's total jobCapacity = Minimum of {CapacityBasedOnMemory, CapacityBasedOnCores}
```

where

```
CapacityBasedOnMemory = (TotalAllocatedMemoryForJobs on Engine / MaxMemory assigned to each Engine Job)
```

```
CapacityBasedOnCores = [Engine's CpuCoreCount - 1]
```

- **Round robin load balancing:** This strategy simply distributes the masking jobs to all the node Continuous Compliance Engines using the round robin algorithm.

### 4.1.1.2 Staging area

The Staging Area is where data from the SOR is unloaded to a series of files by the Hyperscale Compliance Orchestrator. It can be a file system that supports the NFS protocol. The file system can be attached to volumes, or it can be supplied via the Delphix Continuous Data Engine empty VDB feature. In either case, there must be enough storage available to hold the dataset in an uncompressed format. The staging area should be accessible by the Continuous Compliance Engine cluster as well for masking.

### 4.1.1.3 Continuous Compliance Engine cluster

The Continuous Compliance Engine Cluster is a group of Delphix Continuous Compliance Engines (version 6.0.14.0 and later) leveraged by the Hyperscale Compliance Orchestrator to run large masking jobs in parallel. For installing and configuring the Continuous Compliance Engine procedures, see [Continuous Compliance Documentation](#)<sup>20</sup>.

<sup>20</sup> <https://delphixdocs.atlassian.net/wiki/spaces/CC>



#### 4.1.1.4 Source and target data sources

The Hyperscale Compliance Orchestrator is responsible for unloading data from the source data source into a series of files located in the staging area. The Hyperscale Compliance Orchestrator requires network access to the source from the host running the Hyperscale Compliance Orchestrator and credentials to run the appropriate unload commands. After files are masked, the masked data from the files get uploaded to the target data source.

In the case of Oracle and MS SQL data sources, a failure in the load may leave the target data source in an inconsistent state since the load step truncates the target when it begins. If the source and target data source are configured to be the same data source and a failure occurs in the load step, it is recommended that the single data source be restored from a backup (or use the Continuous Data Engine's rewind feature if you have a VDB as the single data source) after the failure in the load step as the data source may be in an inconsistent state. After the data source is restored, you may proceed to kick off another hyperscale job. If the source and target data source are configured to be different, you may use the Hyperscale Compliance Orchestrator restart ability feature to restart the job from the point of failure in the load/post-load step.

##### **Additional data sources:**

- **Delimited Files Connector:** We support hyperscale masking of large delimited files (with a delimiter of single character length). Here the source and target location are considered to be NFS locations.
- **MongoDB Connector:** We support hyperscale masking of large MongoDB database collections. Load step of MongoDB connector drops the target database collection if it is already present. In Hyperscale MongoDB connector, we strongly recommend DONOT use the same collection for both the source and target, particularly when dealing with same MongoDB instances. Utilizing the same collection for in-place masking in such a scenario can pose risks, including potential data deletion, especially when unload, masking, and load operations are occurring asynchronously. It's crucial to maintain a clear separation between source and target entities to ensure data integrity and avoid unintended consequences.
- **Parquet Connector:** We support hyperscale masking of large Parquet files. Here the source and target location are considered to be AWS S3 buckets.



In-Place Masking is NOT supported.

## 4.2 The Continuous Compliance platform

Delphix Continuous Compliance is a multi-user, a browser-based web application that provides complete, secure, and scalable software for your sensitive data discovery, masking, and tokenization needs while meeting enterprise-class infrastructure requirements. To read further about Continuous Compliance features and architecture, read the [Continuous Compliance Documentation](#)<sup>21</sup>.

---

<sup>21</sup> <https://delphixdocs.atlassian.net/wiki/spaces/CC>

## 4.3 Next steps

- Read about [Installation and Setup \(Kubernetes\)](#) (see page 109) .
- Read about the [Network Requirements](#) (see page 73) .
- Read about [Accessing the Hyperscale Compliance API](#) (see page 141) .

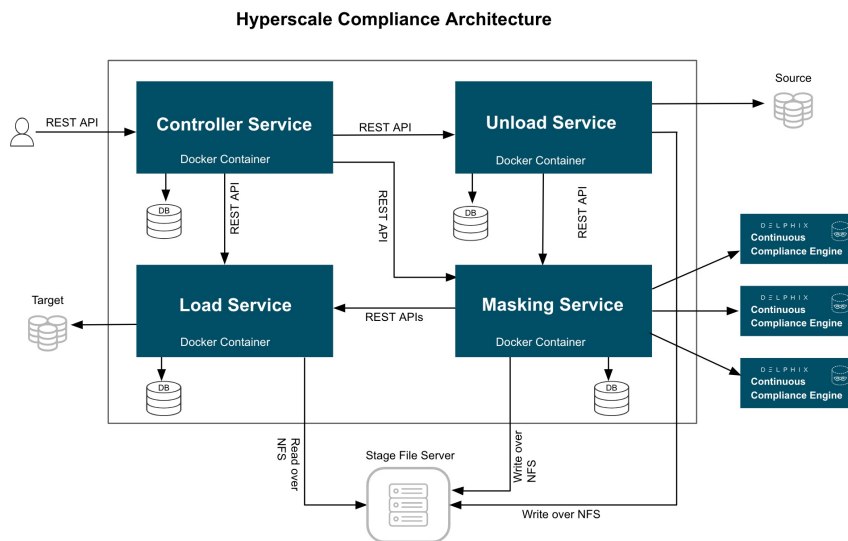
## 5 Getting started

This section covers the following topics:

- [Hyperscale Compliance architecture](#) (see page 59)
- [Data source support](#) (see page 61)
- [Supported platforms](#) (see page 72)
- [Network requirements](#) (see page 73)
- [Deployment](#) (see page 74)
- [NFS server installation](#) (see page 138)
- [Accessing the Hyperscale Compliance API](#) (see page 141)

### 5.1 Hyperscale Compliance architecture

The Hyperscale Compliance architecture comprises four components mainly; Controller Service, Unload Service, Masking Service, and Load Service.



#### 5.1.1 Controller service

The following are the main functions of a controller service:

- Exposes user-accessible API.
- Once the controller service receives user requests (for example, register engine, create a dataset, create a connector, create Job, etc.), it will split the request and sends a request for further processing to downstream services (Unload, Masking, Load) and once response is received from downstream service, the same will be processed by controller service and returned to the user.
- The controller service accepts request job execution from the user and invokes the job execution process by invoking unload service asynchronously.

- The controller service will keep polling data job execution data from the downstream service until execution completes.
- The controller service will also determine the status of job execution and store execution data in the database.
- Controller service allows you to restart a failed (Failed during File Loader, Post Load) execution

### 5.1.2 Unload service

The following are the main functions of a unload service:

- Exposes APIs that are accessible to internal services only.
- Unload service exposes required APIs that help the caller (controller service) to create required inputs (source info, dataset, etc.) for job execution.
- Unload service exposes an API to trigger unload from the source data source. As part of the unload process, it performs the following operations:
  - Reads metadata of source data source (e.g. number of rows in a source file/table) and stores that in the unload service database.
  - Reads data from source data source parallelly (by starting multiple parallel processes for each source entity like tables in case of a relational database ) and stores this data in `.csv` files.
  - Once data is loaded into one `.csv` file, unload service triggers the masking service to start the masking process for that `*.csv` file.
- For running execution, Unload service maintains metadata data (number of rows processed, table/file names processed, etc.) in its database. This data can be retrieved by calling an API.
- Once execution completes execution data in the database and file system gets cleaned by invoking the corresponding API.

### 5.1.3 Masking service

The following are the main functions of a masking service:

- Exposes APIs that are accessible to internal services only.
- Masking services expose required APIs that help the caller (controller service) to create required inputs (Continuous Compliance engine info, dataset, job, etc.) for job execution.
- Masking service exposes an API to trigger the masking process. As part of the masking process, it performs the following operations after receiving a masking request from unload service for a CSV file:
  - Based on Intelligent load balancing, create and start jobs for unloaded files on Continuous Compliance Engines (based on the capacity of Continuous Compliance Engines associated with the hyperscale job).
  - Monitor Continuous Compliance Engine jobs triggered in the previous step.
  - Once monitoring determines that a Continuous Compliance Engine has successfully masked the file, send an async request to the load service (to load data into the target data source) for that masked file.

- For running execution, the Masking service maintains metadata data (number of rows processed, table/file names processed, etc.) in its database. This data can be retrieved by calling an API.
- Once execution completes execution data in the database and file system gets cleaned by invoking the corresponding API.

### 5.1.4 Load service

The following are the main functions of a Load service:

- Exposes APIs that are accessible to internal services only.
- Load service exposes required APIs that help the caller to create required inputs (target data source info, dataset, job, etc.) for job execution.
- Load service exposes an API to trigger the Load process. As part of the Load process, it performs the following operations after receiving a load request from the masking service for a masked CSV file:
  - Perform preload step (for example, cleaning up the target directory or disabling constraints/triggers/indexes). These may be performed once for an execution process (not for each request from the masking service).
  - Load masked files into the target data source.
  - Once Loading for a masked is completed, the metadata for this “file load” will be stored in the load service database.
- For running execution, the Load service maintains metadata data (number of rows processed, table/file names processed, etc.) in its database. This data can be retrieved by calling an API.
- Once execution completes execution data in the database and file system gets cleaned by invoking the corresponding API.
- If the Load service is for a data source that requires post-load steps (e.g. Oracle, MS SQL), then it will include post-load steps which will be triggered by the controller service once all files are successfully loaded into the target data source.
- Load service also allows restarting for the post-load step, if post-load fails for an execution.

## 5.2 Data source support

### 5.2.1 Oracle connector

Oracle Database (commonly referred to as Oracle RDBMS or simply as Oracle) is a multi-model database management system produced and marketed by Oracle Corporation. The following table lists the versions that have been tested in the lab setup:

Platforms	Version
Linux	<ul style="list-style-type: none"> <li>• Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production - AWS</li> <li>• Oracle Database 18c Enterprise Edition Release 18.0.0.0.0 - Production - GCP</li> </ul>



- User on source database must select privileges
- User on target database side must have all privileges and *SELECT\_CATALOG\_ROLE*.

### 5.2.1.1 Supported Data Types

The following are the different data types that are tested in our lab setup:

- VARCHAR
- VARCHAR2
- NUMBER
- FLOAT
- DATE
- TIMESTAMP(default)
- CLOB
- BLOB(with text)
- XMLTYPE



Hyperscale Compliance restricts the support of the following special characters for a database column name: `~!@#%$%^&*()\\"?;:,/\\"`'+=[ ]{|<>'-.\"}]`

### 5.2.1.2 Property values

Property	Value
<code>SKIP.LOAD.SPLIT.COUNT.VALIDATION</code>	false
<code>SKIP.UNLOAD.SPLIT.COUNT.VALIDATION</code>	false

For default values, see [Configuration settings](#) (see page 182).

### 5.2.1.3 Known limitations

The length of the algorithm's generated masked data may exceed the target database table's column length resulting in a job failure. The workaround is to use an algorithm that should generate mask data with a smaller length.

## 5.2.2 MS SQL Connector

### 5.2.2.1 Supported versions

Microsoft SQL Server 2019

### 5.2.2.2 Supported data types

The following are the different data types that are tested in our lab setup:

- VARCHAR
- CHAR
- DATETIME
- INT
- TEXT
- XML (only unload/load)
- VARBINARY (only unload/load)
- SMALLINT
- SMALLMONEY
- MONEY
- BIGINT
- NVARCHAR
- TINYINT
- NUMERIC(X,Y)
- DECIMAL(X,Y)
- FLOAT
- NCHAR
- BIT
- NTEXT
- MONEY

### 5.2.2.3 Property Values

Property	Value
SKIP.LOAD.SPLIT.COUNT.VALIDATION	false
SKIP.UNLOAD.SPLIT.COUNT.VALIDATION	false

For default values, see [Configuration settings \(see page 182\)](#) .

### 5.2.2.4 Known Limitations

- The length of the algorithm's generated masked data may exceed the target database table's column length resulting in a job failure. The workaround is to use an algorithm that should generate mask data with a smaller length.
- The constraints and indexes below are not disabled before load and are enabled back after the load process:
  - Partitioned Indexes

## 5.2.3 Delimited files connector

The connector can be used to mask large delimited files. The delimited unload service splits the large files into smaller chunks and passes them onto the masking service. After the masking is completed, the files are sent to the load service which joins back the split files (the end user also has a choice to disable the join operation).

For Delimited files connector, the splitting/joining of the files is handled by a backend tool i.e. "Data Writer". From the 17.0.0 release and onwards, you can choose the type of "Data Writer" you want to use based on your need as well as understanding the limitations of each type. The supported data writers are:

1. "pyarrow": Apache Arrow is used by the connector to split/join files.
2. "pyspark": Apache Spark is used by the delimited-unload-service to split files. The delimited-load-service will use Linux 'cat' command to join back masked split files.
3. "cat": Only applicable to delimited-load-service, which uses the Linux cat command to join back masked split files.

### 5.2.3.1 Prerequisites

- The source and target (NFS) locations have to be mounted onto the docker containers of unload and load service. Please note that the locations on the containers are what needs to be used when creating the connector-info's using the controller.



```
# As an example
unload-service:
  image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
  ...
  volumes:
    ...
    - /path/to/nfs/mounted/source1/files:/mnt/source1
    - /path/to/nfs/mounted/source2/files:/mnt/source2
  ...
load-service:
  image: delphix-delimited-load-service-app:<HYPERSCALE VERSION>
  ...
  volumes:
    ...
    - /path/to/nfs/mounted/target1/files:/mnt/target1
    - /path/to/nfs/mounted/target2/files:/mnt/target2
```

- Set the required data writer using the `DATA_WRITER_TYPE` environment variable.

```
unload-service:
  image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
  ...
  volumes:
    ...
    - DATA_WRITER_TYPE=pyspark
  ...
load-service:
  image: delphix-delimited-load-service-app:<HYPERSCALE VERSION>
  ...
  environment:
    ...
    - DATA_WRITER_TYPE=pyspark
```

### 5.2.3.2 Property values

Property	Value
<code>SOURCE_KEY_FIELD_NAMES</code>	<code>unique_source_files_identifier</code>
<code>LOAD_SERVICE_REQUIREPOSTLOAD</code>	<code>false</code>

DATA_WRITER_TYPE	<ol style="list-style-type: none"> <li>1. "pyarrow" (Default for delimited-unload-service)</li> <li>2. "pyspark"</li> <li>3. "cat" (Default as well as only applicable to delimited-load-service)</li> </ol>
------------------	--

For default values, see [Configuration settings](#) (see page 182).

### 5.2.3.3 Supported data types

The following are the supported data types for delimited files hyperscale connector:

- String/Text
- Double
- Int64
- Timestamp

### 5.2.3.4 Known limitations

1. Supports only Single-character ASCII delimiters
2. The end-of-record character can only be `\n`, `\r`, or `\r\n`.
3. Limitations with PyArrow Data Writer:
  - a. Output files will exclusively enclose all string types with double quotes ("").
  - b. Columns with double data types will be converted to strings. For example, 6377974237282886994505 will be converted to "36377974237282886994505".
  - c. Columns with int64 data type will be converted to strings. For example, 0009435304391722556805 will be converted to "0009435304391722556805".
4. Limitation with PySpark Data Writer:
  - a. PySpark is more memory intensive, so in case we are processing data that is more in size in comparison to the available memory then we may run into issues related to resource exhaustion. **Caution:** The size of split files multiplied by the number of cores must not exceed the system memory.
  - b. With PyAarrow as the data writer, the split files are generated one after the other, so the masking-service is called as and when a split is created. With PySpark as the data writer, all split files are available only after the split process is complete. So the masking service will be only called after all splits are completed. Due to this, the overall time taken to complete the hyperscale masking execution will be more compared to the former.
  - c. There is a possibility that the number of splits created in the end will be less than the requested number, this generally happens when the file size is small, and spark doesn't create as many partitions as the requested split number.

## 5.2.4 MongoDB connector

The connector can be used to mask large MongoDB files. The Mongo unload service splits the large collections into smaller chunks and passes them onto the masking service. After the masking is completed, the files are sent to the Mongo load service, which imports the masked files into the target collection.

### 5.2.4.1 Supported versions

Platforms	Version
Linux	MongoDB 4.4.x MongoDB 5.0.x MongoDB 6.0.x

### 5.2.4.2 Prerequisites

1. MongoDB users should have the following privileges:

```
use admin
db.createUser({user:"backupadmin", pwd:"xxxxxx", roles:[{role:"backup", db:
"admin"}]})
```

2. Mongo Unload and Mongo Load service image names are to be used under unload-service and load-service. The NFS location has to be mounted onto the Docker containers for unload and load services. Example for mounting `/mnt/hyperscale`.

```
# As an example docker-compose.yaml
unload-service:
  image: delphix-mongo-unload-service-app:${VERSION}
volumes:
  # Uncomment below lines to mount respective paths.
  - /mnt/hyperscale:/etc/hyperscale

load-service:
  image: delphix-mongo-load-service-app:${VERSION}
volumes:
  # Uncomment below lines to mount respective paths.
  - /mnt/hyperscale:/etc/hyperscale
```

2. Uncomment the below lines from `docker-compose.yaml` file under `controller > environment:`

```
# uncomment below for MongoDB connector
#- SOURCE_KEY_FIELD_NAMES=database_name,collection_name
#- VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS=${VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS:-false}
}
#- VALIDATE_MASKED_ROW_COUNT_FOR_STATUS=${VALIDATE_MASKED_ROW_COUNT_FOR_STATUS:-false}
}
#- VALIDATE_LOAD_ROW_COUNT_FOR_STATUS=${VALIDATE_LOAD_ROW_COUNT_FOR_STATUS:-false}
#- DISPLAY_BYTES_INFO_IN_STATUS=${DISPLAY_BYTES_INFO_IN_STATUS:-true}
#- DISPLAY_ROW_COUNT_IN_STATUS=${DISPLAY_ROW_COUNT_IN_STATUS:-false}
```

3. Set the value of `LOAD_SERVICE_REQUIRE_POST_LOAD=false` inside the `“.env”` file.

```
# Set LOAD_SERVICE_REQUIRE_POST_LOAD=false for MongoDB Connector
LOAD_SERVICE_REQUIRE_POST_LOAD=false
```

4. Uncomment the below lines from `“.env”` file.

```
# Uncomment below for MongoDB Connector
#VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS=false
#VALIDATE_MASKED_ROW_COUNT_FOR_STATUS=false
#VALIDATE_LOAD_ROW_COUNT_FOR_STATUS=false
#DISPLAY_BYTES_INFO_IN_STATUS=true
#DISPLAY_ROW_COUNT_IN_STATUS=false
```

### 5.2.4.3 Property values

Mandatory changes are required for the MongoDB Connector in the `docker-compose.yaml` and `.env` files:

Property	Value
SOURCE_KEY_FIELD_NAMES	database_name,collection_name
LOAD_SERVICE_REQUIRE_POST_LOAD	false
VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS	false
VALIDATE_MASKED_ROW_COUNT_FOR_STATUS	false

Property	Value
VALIDATE_LOAD_ROW_COUNT_FOR_STATUS	false
DISPLAY_BYTES_INFO_IN_STATUS	true
DISPLAY_ROW_COUNT_IN_STATUS	false

For default values, see [Configuration settings](#)<sup>22</sup>.

#### 5.2.4.4 Known limitation:

- Sharded MongoDB Atlas is not supported.
- In-Place Masking is not supported.

### 5.2.5 Parquet connector

The connector can be used to mask large Parquet files. The parquet unload service splits the large files into smaller chunks and passes them onto the masking service. After the masking is completed, the files are sent to the load service, which joins back the split files (you also have a choice to disable the join operation).

#### 5.2.5.1 Prerequisites

- The connector should be able to access the AWS S3 buckets (the source and target locations). The following approaches are supported by the connector and can be used to authenticate with the S3 bucket:
  - Attaching the IAM role to the EC2 instance where the hyperscale masking services will be deployed.
    - IAM Roles are designed for applications to securely make AWS-API requests from EC2 instances, without the necessity to manage the security credentials that the applications use.
    - Using the AWS console UI or AWS CLI, attach the IAM role to the EC2 instance running the Hyperscale services. To know more, check the [AWS Documentation](#)<sup>23</sup>.
    - With IAM role authentication, there is no need to pass the AWS credentials during the connector-info creation.

```
# Example connector-info payload
```

<sup>22</sup> <https://hyperscalemasking.delphix.com/docs/latest/configuration-settings>

<sup>23</sup> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

```

{
  "source": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "aws_s3_bucket/sub_folder(s)"
    }
  },
  "target": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "aws_s3_bucket/sub_folder(s)"
    }
  }
}

```

- Passing the AWS Access Key ID & AWS Secret Access Key attached to an AWS role:
  - Access keys are long-term credentials generated for an IAM user or role. These keys can be for programmatic requests to the AWS CLI or AWS API (directly or using the AWS SDK). To know more, check the [AWS Documentation](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)<sup>24</sup>.
  - These credentials can be passed during the connector-info creation.

```

# Example connector-info payload
{
  "source": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "aws_s3_bucket/sub_folder(s)",
      "aws_region": "us-west-2",
      "aws_access_key_id": "AWS_ACCESS_KEY_ID",
      "aws_secret_access_key": "AWS_SECRET_ACCESS_KEY"
    }
  },
  "target": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "aws_s3_bucket/sub_folder(s)",
      "aws_region": "us-west-2",
      "aws_access_key_id": "AWS_ACCESS_KEY_ID",
      "aws_secret_access_key": "AWS_SECRET_ACCESS_KEY"
    }
  }
}

```

<sup>24</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_access-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)

- They can also be set as environment variables when bringing up the Parquet connector services.

```

unload-service:
  ...
  environment:
    - AWS_DEFAULT_REGION=us-east-1
    - AWS_ACCESS_KEY_ID=<aws_access_key_id>
    - AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>
  ...
load-service:
  ...
  environment:
    - AWS_DEFAULT_REGION=us-east-1
    - AWS_ACCESS_KEY_ID=<aws_access_key_id>
    - AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>

```

### 5.2.5.2 Property values

Configurations on the controller service:

Property	Value
SOURCE_KEY_FIELD_NAMES	unique_source_files_identifier
LOAD_SERVICE_REQUIREPOSTLOAD	false

Configuration on the parquet-unload-service:

Property	Value
MAX_WORKER_THREADS_PER_JOB	512

For default values, see [Configuration settings](#)<sup>25</sup>.

### 5.2.5.3 Supported data types

The following are the supported data types for parquet files hyperscale connector:

<sup>25</sup> <https://hyperscalemasking.delphix.com/docs/latest/configuration-settings>

- BOOLEAN
- INT32
- INT64
- INT96
- FLOAT
- DOUBLE
- BYTE\_ARRAY

#### 5.2.5.4 Known limitations


- Generally, the parquet files are compressed and the compression factor could vary from 2x to 70x or even more. So, when working with such larger files the connector will need a host which has large enough memory to accommodate the parallel execution of multiple large parquet files. ***In case the sum of the uncompressed size of parquet files that are getting executed in parallel exceeds 80% of RAM size then the chances of having an “out of memory” error are high.*** To avoid OOM, the end user can reduce the MAX\_WORKER\_THREADS\_PER\_JOB (i.e. reduce the number of parallel threads), ultimately reducing the memory usage.

### 5.3 Supported platforms

Delphix supports Hyperscale Compliance for many data platforms and operating systems.

#### 5.3.1 Supported Continuous Compliance/Data versions

Delphix Engine	Minimum supported version	Recommended version
Continuous Compliance	6.0.14.0 15.0.0.0 (for embedded XML/JSON masking)	Latest
Continuous Data	6.0.14.0	Latest

 All Continuous Compliance Engines must be of the same versions and must be used only by Hyperscale Compliance for masking. Already existing or running masking/profiling jobs on Continuous Compliance engines would impact Hyperscale Compliance performance and results.



### 5.3.2 Supported browsers (only API client)

Hyperscale Compliance API Client is using Swagger UI-3.48.0 which works in the latest versions of Chrome, Safari, Firefox, and Edge. For more information about the supported browser versions, see the **Browser Support** section on [GitHub](#)<sup>26</sup>.

- i** If you encounter Chrome `NET::ERR_CERT_INVALID` error code, perform the following steps to resolve the above error:
- Type `https://<hyperscale-compliance-host address>/hyperscale-compliance` in the address bar and click **Enter**.
  - Right-click on the page and click **Inspect**.
  - Click the **Console** tab and run the following command:  
`sendCommand(SecurityInterstitialCommandId.CMD_PROCEED)`.
  - Click on **Authorize** and provide the key. For more information about the key, refer to step 7 in [Generate a New Key](#) (see page 76).

## 5.4 Network requirements

This section describes the network requirements for Hyperscale Compliance. Ensure that you meet all the network requirements before you install the Hyperscale Compliance Orchestrator.

The following are the inbound/outbound rules for the Hyperscale Compliance Orchestrator:

Type (Inbound/Outbound)	Port	Reason
Inbound and Outbound	80	HTTP connections to/from the Hyperscale Compliance Orchestrator to/from the Continuous Compliance Engines part of the Continuous Compliance Engine Cluster and to access the Hyperscale Compliance API.

<sup>26</sup> <https://github.com/swagger-api/swagger-ui#browser-support>

Type (Inbound/Outbound)	Port	Reason
Inbound and Outbound	443	HTTPs connections to/from the Hyperscale Compliance Orchestrator to/from the Continuous Compliance Engines part of the Continuous Compliance Engine Cluster and to access the Hyperscale Compliance API.
Outbound	53	Connections to local DNS servers.
Inbound	22	SSH connections to the Hyperscale Compliance Orchestrator host.

## 5.5 Deployment

This section covers the following topics:

- [Docker compose](#) (see page 74)
- [Kubernetes](#) (see page 106)
- [Podman compose](#) (see page 119)

### 5.5.1 Docker compose



Delphix has announced the deprecation of support for Docker Compose with Hyperscale version 17.0.0. 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. It is highly recommended that new Hyperscale installations be performed on Kubernetes.

This section covers the following topics:

- [Host requirements \(Docker compose\)](#) (see page 75)
- [Installation and setup \(Docker compose\)](#) (see page 76)
- [Custom configuration](#) (see page 94)
- [Upgrading the Hyperscale Compliance orchestrator \(Docker compose\)](#) (see page 95)
- [How to generate a support bundle \(Docker compose\)](#) (see page 98)
- [Managing the storage space](#) (see page 100)

- [Migrating to Kubernetes \(see page 101\)](#)

### 5.5.1.1 Host requirements (Docker compose)

Type	Host Requirement	Explanation
User	<p>A user (hyperscale_os) with the following permissions are required:</p> <ul style="list-style-type: none"> <li>• Should have permissions to install <code>docker</code> and <code>docker-compose</code>.</li> <li>• Should be part of the 'docker' OS group or must have the permission to run <code>docker</code> and <code>docker-compose</code> commands.</li> <li>• Permission to run <code>mount</code>, <code>unmount</code>, <code>mkdir</code> and <code>rmdir</code> as a super-user with <code>NOPASSWD</code>.</li> <li>• Should have either <code>GID=50</code> and/or <code>UID=65436</code>.</li> </ul>	<p>This will be a primary user responsible to install and operate the Hyperscale Compliance.</p>
Installation Directory	<p>There must be a directory on the Hyperscale Compliance Orchestrator host where the Hyperscale Compliance can be installed.</p>	<p>This is a directory where the Hyperscale Compliance tar archive file will be placed and extracted. The extracted artifacts will include docker images (tar archive files) and a configuration file (<code>docker-compose.yaml</code>) that will be used to install the Hyperscale Compliance.</p>
Log File Directory	<p>An optional directory to place log files.</p>	<p>This directory (can be configured via <code>docker-compose.yaml</code> configuration file) will host the runtime/log files of the Hyperscale Compliance Orchestrator.</p>

Type	Host Requirement	Explanation
NFS Client Services	NFS client services must be enabled on the host.	NFS client service is required to be able to mount an NFS shared storage from where the Hyperscale Compliance Orchestrator will be able to read the source files and write the target files. For more information, see <a href="#">NFS Server Installation</a> (see page 138).
Hardware Requirements	<ul style="list-style-type: none"> <li>• Minimum: 8 vCPU, 64 GB of memory, 100GB data disk.</li> <li>• Recommended: 16 vCPU, 128GB of memory, 500GB data disk.</li> </ul>	OS disk space: 50 GB

### 5.5.1.2 Installation and setup (Docker compose)



Delphix has announced the depreciation of support for Docker Compose with Hyperscale version 17.0.0. 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. It is highly recommended that new Hyperscale installations be performed on Kubernetes.

This section describes the steps you must perform to install the Hyperscale Compliance Orchestrator.

#### 5.5.1.2.1 Hyperscale Compliance installation


##### 5.5.1.2.1.1 Pre-requisites

Ensure that you meet the following requirements before you install the Hyperscale Compliance Orchestrator.

- Download the Hyperscale tar file ( `delphix-hyperscale-masking-17.0.0.tar.gz` ) from [download.delphix.com](https://download.delphix.com)<sup>27</sup>.
- You must create a user that has permission to install Docker and Docker Compose.
- Install Docker on VM. The minimum supported docker version is 20.10.7.

<sup>27</sup> <https://download.delphix.com/folder/2224/Delphix%20Product%20Releases/Hyperscale%20Compliance>

- Install Docker Compose on the VM. The minimum supported docker-compose version is 1.29.2.
- Check if docker and docker-compose are installed by running the following command:
  - `docker-compose -v` The above command displays an output similar to the following:  
`docker-compose version 1.29.2, build 5becea4c`
  - `docker -v` The above command displays an output similar to the following: `Docker version 20.10.7, build 3967b7d`
- [Only Required for Oracle Load Service] Download and install Linux-based [Oracle's instant client](#)<sup>28</sup> on the machine where the Hyperscale Compliance Orchestrator will be installed. The client should essentially include `instantclient-basic` (Oracle shared libraries) along with `instantclient-tools` containing Oracle's SQL\*Loader client. Both the packages `instantclient-basic` and `instantclient-tools` should be unzipped in the same directory. A group ownership id of 50 with a permission mode of 550 or a user id of 65436 with a permission mode of 500 must be set recursively on the directory where Oracle's instant client `binaries/libraries` will be installed. This is required by the Hyperscale Compliance Orchestrator to be able to read or execute from the directory.

 Oracle Load doesn't support Object Identifiers(OIDs).

#### 5.5.1.2.1.2 Procedure

Perform the following procedure to install the Hyperscale Compliance Orchestrator.

1. Unpack the Hyperscale tar file.

```
tar -xzf delphix-hyperscale-masking-17.0.0.tar.gz
```

2. Upon unpacking, you will find the docker image tar files which are categorized as below:

##### **Universal images common for all connectors.**

- `controller-service.tar`
- `masking-service.tar`
- `proxy.tar`

##### **Oracle (required only for Oracle data source masking)**

- `unload-service.tar`
- `load-service.tar`

##### **MSSQL (required only for MS SQL data source masking)**

- `mssql-unload-service.tar`

<sup>28</sup> <https://www.oracle.com/database/technologies/instant-client.html>

- `mssql-load-service.tar`

#### **Delimited Files (required only for Delimited Files masking)**

- `delimited-unload-service.tar`
- `delimited-load-service.tar`

#### **MongoDB (required only for MongoDB database masking)**

- `mongo-unload-service.tar`
- `mongo-load-service.tar`

#### **Parquet files (required only for Parquet file masking)**

- `parquet-unload-service.tar`
- `parquet-load-service.tar`

Each deployment set consists of 5 images (3 Universal images and 2 images related to each dataset type). Proceed to load the required images into Docker as below:

For Oracle data source masking:

```
docker load --input unload-service.tar
docker load --input load-service.tar
docker load --input controller-service.tar
docker load --input masking-service.tar
docker load --input proxy.tar
```

For MS SQL data source masking:

```
docker load --input mssql-unload-service.tar
docker load --input mssql-load-service.tar
docker load --input controller-service.tar
docker load --input masking-service.tar
docker load --input proxy.tar
```

For Delimited Files masking:

```
docker load --input delimited-unload-service.tar
docker load --input delimited-load-service.tar
docker load --input controller-service.tar
docker load --input masking-service.tar
docker load --input proxy.tar
```

For MongoDB data source masking:

```
docker load --input mongo-unload-service.tar
docker load --input mongo-load-service.tar
```

```
docker load --input controller-service.tar
docker load --input masking-service.tar
docker load --input proxy.tar
```

For Parquet masking:

```
docker load --input parquet-unload-service.tar
docker load --input parquet-load-service.tar
docker load --input controller-service.tar
docker load --input masking-service.tar
docker load --input proxy.tar
```

3. Create an NFS shared mount, that will act as a **Staging Area**, on the Hyperscale Compliance Orchestrator host where the Hyperscale Compliance Orchestrator will perform read/write/execute operations:

1. Create a 'Staging Area' directory. For example: `/mnt/hyperscale/staging_area`. The user(s) within each of the docker containers part of the Hyperscale Compliance Orchestrator and the appliance OS user(s) in the Continuous Compliance Engine(s), all have the user id as 65436 and/or group ownership id as 50. As such, the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, require the following permissions, based on the UID/GID of the OS user, so that the Hyperscale Compliance Orchestrator and the Continuous Compliance Engine(s) can perform read/write/execute operations on the staging area:
  - If the Hyperscale Compliance OS user has a UID of 65436, then the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, must have a UID of 65436 and 700 permission mode.
  - If the Hyperscale Compliance OS user has a GID of 50 and does not have a UID of 65436, then the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, must have a GID of 50 and 770 permission mode.
2. Mount the NFS shared directory on the staging area directory( `/mnt/hyperscale/staging_area` ). This NFS shared storage can be created and mounted in two ways as detailed in the [NFS Server Installation](#)<sup>29</sup> section. Based on the umask value for the user which is used to mount, the permissions for the staging area directory could get altered after the NFS share has been mounted. In such cases, the permissions(i.e. 770 or 700 whichever applies based on point 3a) must be applied again on the staging area directory.

**i** The directory created in step 3a ('staging\_area') will be provided as the `mountName` and the corresponding shared path from the NFS file server as the `mountPath` in the MountFileSystems API.

<sup>29</sup> <https://portal.document360.io/continuous-compliance/docs/nfs-server-installation>

4. After unpacking the tar, you will find the following sample docker-compose files available, `docker-compose-sample.yaml`, `docker-compose-oracle-sample.yaml`, `docker-compose-mssql-sample.yaml`, `docker-compose-delimitedfiles-sample.yaml`, `docker-compose-mongo-sample.yaml`, and `docker-compose-parquet-sample.yaml`. These sample files can be used to create a `docker-compose.yaml` file based on the connector you want to use. Configure the following docker container volume bindings for the docker containers by editing the `docker-compose.yaml`:

1. For each of the docker containers, except the 'proxy' container, add a volume entry binding the staging area path (from 3(a), `/mnt/hyperscale`) to the Hyperscale Compliance Orchestrator container path (`/etc/hyperscale`) as a volume binding under the 'volumes' section.
2. [Only Required for Oracle Load Service] For the **load-service** docker container, add a volume entry that binds the path of the directory on the host where both the Oracle instant Client packages were unzipped to the path on the container (`/usr/lib/instantclient`) under the 'volumes' section.
3. [Only Required for Delimited Unload Service] For Delimited Files unload-service, the source NFS location has to be mounted to the container as docker volume in order for it to access the source files. The path mounted on the container is passed during the creation of the source connector-info.

```
# Mount your source NFS location onto your Hyperscale Engine server
sudo mount [-t nfs4] <source_nfs_endpoint>:<source_nfs_location>
<nfs_mount_on_host>
```

```
# Later mount <nfs_mount_on_host> as a docker volume to the delimited unload-
service container (in docker-compose.yaml, created using docker-compose-
delimitedfiles-sample.yaml)
unload-service:
  image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
  ...
  volumes:
    ...
    # Source files should be made available within the unload-service
container file system
    # The paths within the container should be configured in the source
section of connector-info [with type=FS]
    -
    <nfs_mount_on_host>:<source_files_mount_passed_during_connector_info_creation_p
ath_in_contianer>
```

4. [Only Required for Delimited load Service] For Delimited Files load-service, the target NFS location has to be mounted to the container as docker volume in order for it to access the target location where masked files will be placed. The path mounted on the container is passed during the creation of the target connector-info.



```
# Mount your target NFS location onto your Hyperscale Engine server
sudo mount [-t nfs4] <target_nfs_endpoint>:<target_nfs_location>
<target_nfs_mount_on_host>
```

```
# Later mount <nfs_mount_on_host> as a docker volume to the delimited load-
service container (in docker-compose.yaml, created using docker-compose-
delimitedfiles-sample.yaml)
load-service:
  image: delphix-delimited-load-service-app:${VERSION}
  ...
  volumes:
    ...
    # Target location should be made available within the load-service
container file system
    # The paths within the container should be configured in the target
section of connector-info [with type=FS]
    -
    <target_nfs_mount_on_host>:<target_location_passed_during_connector_info_creati
on_in_container>
```

5. [Optional] Some data (for example, logs, configuration files, etc.) that is generated inside the docker containers may be useful to debug possible errors or exceptions while running the hyperscale jobs, and as such it may be beneficial to persist these logs outside docker containers. The following data can be persisted outside the docker containers:
- The logs generated for each service i.e. unload, controller, masking, and load services.
  - The sqldr utility logs and control files at `opt/sqldr` location in the load-service container.
  - The file-upload folder at `/opt/delphix/uploads` in the controller-service container

If you would like to persist the above data on your host, then you have the option to do the same by setting up volume bindings in the respective service as indicated below, that map locations inside the docker containers to locations on the host in the `docker-compose.yaml` file. The host locations again must have a group ownership id of 50 with a permission mode of 770 or a user id of 65436 with a permission of 700, due to the same reasons as highlighted in step 3a.

Here are examples of the `docker-compose.yaml` file for Oracle, MS SQL, MongoDB data sources, Delimited file masking, and Parquet file masking:

For Oracle data source masking:

```
version: "3.7"
services:
  controller-service:
    image: delphix-controller-service-app:${VERSION}
    healthcheck:
      test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
      interval: 30s
      timeout: 25s
```

```

    retries: 3
    start_period: 30s
depends_on:
  - unload-service
  - masking-service
  - load-service
init: true
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-controller-data:/data
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /home/hyperscale_user/logs/controller_service:/opt/delphix/logs
  - /mnt/hyperscale:/etc/hyperscale
environment:
  - API_KEY_CREATE=${API_KEY_CREATE:-false}
  - EXECUTION_STATUS_POLL_DURATION=${EXECUTION_STATUS_POLL_DURATION:-12000}
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_CONTROLLER_SERVICE:-INFO}
  - API_VERSION_COMPATIBILITY_STRICT_CHECK=${
{API_VERSION_COMPATIBILITY_STRICT_CHECK:-false}
  - LOAD_SERVICE_REQUIREPOSTLOAD=${LOAD_SERVICE_REQUIRE_POST_LOAD:-true}
  - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_UNLOAD_SPLIT_COUNT_VALIDATION:-false}
  - SKIP_LOAD_SPLIT_COUNT_VALIDATION=${SKIP_LOAD_SPLIT_COUNT_VALIDATION:-false}
  - CANCEL_STATUS_POLL_DURATION=${CANCEL_STATUS_POLL_DURATION:-60000}
unload-service:
image: delphix-unload-service-app:${VERSION}
init: true
environment:
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_UNLOAD_SERVICE:-INFO}
  - UNLOAD_FETCH_ROWS=${UNLOAD_FETCH_ROWS:-10000}
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-unload-data:/data
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/hyperscale:/etc/hyperscale
  - /home/hyperscale_user/logs/unload_service:/opt/delphix/logs
masking-service:
image: delphix-masking-service-app:${VERSION}
init: true
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-masking-data:/data
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/hyperscale:/etc/hyperscale

```

```

- /home/hyperscale_user/logs/masking_service:/opt/delphix/logs
environment:
- LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_MASKING_SERVICE:-INFO}
- INTELLIGENT_LOADBALANCE_ENABLED=${INTELLIGENT_LOADBALANCE_ENABLED:-true}
load-service:
image: delphix-load-service-app:${VERSION}
init: true
environment:
- LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_LOAD_SERVICE:-INFO}
- SQLLDR_BLOB_CLOB_CHAR_LENGTH=${SQLLDR_BLOB_CLOB_CHAR_LENGTH:-20000}
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-load-data:/data
# The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
- /mnt/hyperscale:/etc/hyperscale
- /opt/oracle/instantclient_21_5:/usr/lib/instantclient
- /home/hyperscale_user/logs/load_service:/opt/delphix/logs
- /home/hyperscale_user/logs/load_service/sqlldr:/opt/sqlldr/
proxy:
image: delphix-hyperscale-masking-proxy:${VERSION}
init: true
networks:
- hyperscale-net
ports:
- "443:443"
restart: unless-stopped
depends_on:
- controller-service
#volumes:
# Uncomment to bind mount /etc/config
#- /nginx/config/path/on/host:/etc/config
networks:
hyperscale-net:
volumes:
hyperscale-load-data:
hyperscale-unload-data:
hyperscale-masking-data:
hyperscale-controller-data:

```

For MS SQL data source masking:

```

version: "3.7"
services:
controller-service:
image: delphix-controller-service-app:${VERSION}
healthcheck:
test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
interval: 30s

```

```

    timeout: 25s
    retries: 3
    start_period: 30s
depends_on:
  - unload-service
  - masking-service
  - load-service
init: true
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-controller-data:/data
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /home/hyperscale_user/logs/controller_service:/opt/delphix/logs
  - /mnt/hyperscale:/etc/hyperscale
environment:
  - API_KEY_CREATE=${API_KEY_CREATE:-false}
  - EXECUTION_STATUS_POLL_DURATION=${EXECUTION_STATUS_POLL_DURATION:-12000}
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_CONTROLLER_SERVICE:-INFO}
  - API_VERSION_COMPATIBILITY_STRICT_CHECK=${
{API_VERSION_COMPATIBILITY_STRICT_CHECK:-false}
  - LOAD_SERVICE_REQUIREPOSTLOAD=${LOAD_SERVICE_REQUIRE_POST_LOAD:-true}
  - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_UNLOAD_SPLIT_COUNT_VALIDATION:-false}
  - SKIP_LOAD_SPLIT_COUNT_VALIDATION=${SKIP_LOAD_SPLIT_COUNT_VALIDATION:-false}
  - CANCEL_STATUS_POLL_DURATION=${CANCEL_STATUS_POLL_DURATION:-60000}
unload-service:
  image: delphix-mssql-unload-service-app:${VERSION}
  init: true
  environment:
    - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_UNLOAD_SERVICE:-INFO}
    - UNLOAD_FETCH_ROWS=${UNLOAD_FETCH_ROWS:-10000}
    - SPARK_DATE_TIMESTAMP_FORMAT=${DATE_TIMESTAMP_FORMAT:-yyyy-MM-dd
HH:mm:ss.SSS}
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-unload-data:/data
    # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
    - /mnt/hyperscale:/etc/hyperscale
    - /home/hyperscale_user/logs/unload_service:/opt/delphix/logs
masking-service:
  image: delphix-masking-service-app:${VERSION}
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-masking-data:/data

```

```

# The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
- /mnt/hyperscale:/etc/hyperscale
- /home/hyperscale_user/logs/masking_service:/opt/delphix/logs
environment:
- LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_MASKING_SERVICE:-INFO}
- INTELLIGENT_LOADBALANCE_ENABLED=${INTELLIGENT_LOADBALANCE_ENABLED:-true}
load-service:
image: delphix-mssql-load-service-app:${VERSION}
init: true
environment:
- LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_LOAD_SERVICE:-INFO}
- SQLLDR_BLOB_CLOB_CHAR_LENGTH=${SQLLDR_BLOB_CLOB_CHAR_LENGTH:-20000}
- SPARK_DATE_TIMESTAMP_FORMAT=${DATE_TIMESTAMP_FORMAT:-yyyy-MM-dd
HH:mm:ss.SSSS}
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-load-data:/data
# The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
- /mnt/hyperscale:/etc/hyperscale
- /home/hyperscale_user/logs/load_service:/opt/delphix/logs
proxy:
image: delphix-hyperscale-masking-proxy:${VERSION}
init: true
networks:
- hyperscale-net
ports:
- "443:443"
restart: unless-stopped
depends_on:
- controller-service
#volumes:
# Uncomment to bind mount /etc/config
#- /nginx/config/path/on/host:/etc/config
networks:
hyperscale-net:
volumes:
hyperscale-load-data:
hyperscale-unload-data:
hyperscale-masking-data:
hyperscale-controller-data:

```

For Delimited Files masking:

A sample file specific to the Delimited connector is available in the package called `docker-compose-delimitedfiles-sample.yaml`.

```

version: "3.7"
services:

```

```

controller-service:
  image: delphix-controller-service-app:<HYPERSCALE VERSION>
  healthcheck:
    test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
    interval: 30s
    timeout: 25s
    retries: 3
    start_period: 30s
  depends_on:
    - unload-service
    - masking-service
    - load-service
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-controller-data:/data
    # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
    - /mnt/parent_staging_area:/etc/hyperscale
  environment:
    - API_KEY_CREATE=true
    - EXECUTION_STATUS_POLL_DURATION=120000
    - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
    - API_VERSION_COMPATIBILITY_STRICT_CHECK=false
    - LOAD_SERVICE_REQUIREPOSTLOAD=false
    - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=false
    - SKIP_LOAD_SPLIT_COUNT_VALIDATION=false
    - CANCEL_STATUS_POLL_DURATION=60000
    - SOURCE_KEY_FIELD_NAMES=unique_source_files_identifier
unload-service:
  image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-unload-data:/data
    # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
    # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
    - /mnt/parent_staging_area:/etc/hyperscale
    # Source files should be made available within the unload-service container
file system
    # The paths within the container should be configured in the source section of
connector-info [with type=FS]
    - /mnt/source_files:/mnt/source
    #- /mnt/source_files2:/mnt/source2
masking-service:
  image: delphix-masking-service-app:<HYPERSCALE VERSION>

```

```

init: true
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-masking-data:/data
  # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/parent_staging_area:/etc/hyperscale
environment:
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
  - INTELLIGENT_LOADBALANCE_ENABLED=true
load-service:
  image: delphix-delimited-load-service-app:<HYPERSCALE VERSION>
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-load-data:/data
    # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
    # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
    - /mnt/parent_staging_area:/etc/hyperscale
    # Target location should be made available within the load-service container
file system
    # The paths within the container should be configured in the target section of
connector-info [with type=FS]
    - /mnt/target_files:/mnt/target
    #- /mnt/target_files2:/mnt/target2
proxy:
  image: delphix-hyperscale-masking-proxy:<HYPERSCALE VERSION>
  init: true
  networks:
    - hyperscale-net
  ports:
    - "443:443"
    - "80:80"
  restart: unless-stopped
  depends_on:
    - controller-service
networks:
  hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data:

```

For MongoDB data source masking:

```

# Copyright (c) 2021, 2023 by Delphix. All rights reserved.
version: "3.7"
services:
  controller-service:
    build:
      context: controller-service
      args:
        - VERSION=${VERSION}
    image: delphix-controller-service-app:${VERSION}
    healthcheck:
      test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP
|| exit 1'
      interval: 30s
      timeout: 25s
      retries: 3
      start_period: 30s
    depends_on:
      - unload-service
      - masking-service
      - load-service
    init: true
    networks:
      - hyperscale-net
    restart: unless-stopped
    volumes:
      - hyperscale-controller-data:/data
      # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
      - /home/hyperscale_user/logs/controller_service:/opt/delphix/logs
      - /mnt/hyperscale:/etc/hyperscale

    environment:
      - API_KEY_CREATE=${API_KEY_CREATE:-false}
      - EXECUTION_STATUS_POLL_DURATION=${EXECUTION_STATUS_POLL_DURATION:-120000}
      - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_CONTROLLER_SERVICE:-
INFO}
      - API_VERSION_COMPATIBILITY_STRICT_CHECK=${
{API_VERSION_COMPATIBILITY_STRICT_CHECK:-false}
      - LOAD_SERVICE_REQUIREPOSTLOAD=${LOAD_SERVICE_REQUIRE_POST_LOAD:-true}
      - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_UNLOAD_SPLIT_COUNT_VALIDATION:-false}
      - SKIP_LOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_LOAD_SPLIT_COUNT_VALIDATION:-false}
      - CANCEL_STATUS_POLL_DURATION=${CANCEL_STATUS_POLL_DURATION:-60000}
      - SOURCE_KEY_FIELD_NAMES=database_name,collection_name
      - VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS=${
{VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS:-false}
      - VALIDATE_MASKED_ROW_COUNT_FOR_STATUS=${
{VALIDATE_MASKED_ROW_COUNT_FOR_STATUS:-false}
      - VALIDATE_LOAD_ROW_COUNT_FOR_STATUS=${
{VALIDATE_LOAD_ROW_COUNT_FOR_STATUS:-false}
      - DISPLAY_BYTES_INFO_IN_STATUS=${DISPLAY_BYTES_INFO_IN_STATUS:-true}

```



```

- DISPLAY_ROW_COUNT_IN_STATUS=${DISPLAY_ROW_COUNT_IN_STATUS:-false}

unload-service:
  build:
    context: unload-service
    args:
      - VERSION=${VERSION}
  image: delphix-mongo-unload-service-app:${VERSION}
  init: true
  environment:
    - LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_UNLOAD_SERVICE:-INFO}
    - UNLOAD_FETCH_ROWS=${UNLOAD_FETCH_ROWS:-10000}
    - CONCURRENT_EXPORT_LIMIT=${CONCURRENT_EXPORT_LIMIT:-10}
    - HIKARI_MAX_LIFE_TIME=${UNLOAD_HIKARI_MAX_LIFE_TIME:-1800000}
    - HIKARI_KEEP_ALIVE_TIME=${UNLOAD_HIKARI_KEEP_ALIVE_TIME:-300000}
    - FILE_DELIMITER=${FILE_DELIMITER:-,}
    - FILE_ENCLOSURE=${FILE_ENCLOSURE:-"}
    - FILE_ESCAPE_ENCLOSURE=${FILE_ESCAPE_ENCLOSURE:-"}

  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-unload-data:/data
    # The orchestrator VM paths(left side of colon) used here are examples.
    Configure the respective mount paths.
    - /mnt/hyperscale:/etc/hyperscale
    - /home/hyperscale_user/logs/unload_service:/opt/delphix/logs

masking-service:
  build:
    context: masking-service
    args:
      - VERSION=${VERSION}
  image: delphix-masking-service-app:${VERSION}
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-masking-data:/data
    # The orchestrator VM paths(left side of colon) used here are examples.
    Configure the respective mount paths.
    - /mnt/hyperscale:/etc/hyperscale
    - /home/hyperscale_user/logs/masking_service:/opt/delphix/logs
  environment:
    - LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_MASKING_SERVICE:-INFO}
    - INTELLIGENT_LOADBALANCE_ENABLED=${INTELLIGENT_LOADBALANCE_ENABLED:-true}

load-service:
  build:
    context: load-service
    args:

```

```

    - VERSION=${VERSION}
image: delphix-mongo-load-service-app:${VERSION}
init: true
environment:
  - LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_LOAD_SERVICE:-INFO}
  - SQLLDR_BLOB_CLOB_CHAR_LENGTH=${SQLLDR_BLOB_CLOB_CHAR_LENGTH:-20000}
  - HIKARI_MAX_LIFE_TIME=${LOAD_HIKARI_MAX_LIFE_TIME:-1800000}
  - HIKARI_KEEP_ALIVE_TIME=${LOAD_HIKARI_KEEP_ALIVE_TIME:-300000}
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-load-data:/data
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/hyperscale:/etc/hyperscale
  - /home/hyperscale_user/logs/load_service:/opt/delphix/logs

proxy:
  build: nginx
  image: delphix-hyperscale-masking-proxy:${VERSION}
  init: true
  networks:
    - hyperscale-net
  ports:
    - "443:443"
    - "80:80"
  restart: unless-stopped
  depends_on:
    - controller-service
  #volumes:
    # Uncomment to bind mount /etc/config
    #- /nginx/config/path/on/host:/etc/config
networks:
  hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data:

```

For Parquet files masking:

A sample file specific to the Parquet connector is available in the package called `docker-compose-parquet-sample.yaml`.

```

version: "3.7"
services:
  controller-service:
    image: delphix-controller-service-app:<HYPERSCALE VERSION>
    healthcheck:

```

```

test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
interval: 30s
timeout: 25s
retries: 3
start_period: 30s
depends_on:
- unload-service
- masking-service
- load-service
init: true
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-controller-data:/data
# The orchestrator VM paths(left side of colon) used here are examples.

```

Configure the respective mount paths.

```
- /mnt/parent_staging_area:/etc/hyperscale
```

environment:

```

- API_KEY_CREATE=true
- EXECUTION_STATUS_POLL_DURATION=120000
- LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=INFO
- API_VERSION_COMPATIBILITY_STRICT_CHECK=false
- LOAD_SERVICE_REQUIREPOSTLOAD=false
- SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=false
- SKIP_LOAD_SPLIT_COUNT_VALIDATION=false
- CANCEL_STATUS_POLL_DURATION=60000
- SOURCE_KEY_FIELD_NAMES=unique_source_files_identifier

```

unload-service:

```
image: delphix-parquet-unload-service-app:<HYPERSCALE VERSION>
```

```
init: true
```

networks:

```
- hyperscale-net
```

```
restart: unless-stopped
```

volumes:

```

- hyperscale-unload-data:/data
# Staging area volume mount, here /mnt/parent_staging_area is used as an example
# The orchestrator VM paths(left side of colon) used here are examples.

```

Configure the respective mount paths.

```
- /mnt/parent_staging_area:/etc/hyperscale
```

environment:

```

- MAX_WORKER_THREADS_PER_JOB=512
# The default AWS region and credentials can be set using environment variables
#- AWS_DEFAULT_REGION=us-east-1
#- AWS_ACCESS_KEY_ID=<aws_access_key_id>
#- AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>

```

masking-service:

```
image: delphix-masking-service-app:<HYPERSCALE VERSION>
```

```
init: true
```

networks:

```
- hyperscale-net
```

```
restart: unless-stopped
```

```

volumes:
  - hyperscale-masking-data:/data
  # Staging area volume mount, here /mnt/parent_staging_area is used as an example
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/parent_staging_area:/etc/hyperscale
environment:
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
  - INTELLIGENT_LOADBALANCE_ENABLED=true
load-service:
image: delphix-parquet-load-service-app:<HYPERSCALE VERSION>
init: true
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-load-data:/data
  # Staging area volume mount, here /mnt/parent_staging_area is used as an example
  # The orchestrator VM paths(left side of colon) used here are examples.
Configure the respective mount paths.
  - /mnt/parent_staging_area:/etc/hyperscale
#environment:
# The default AWS region and credentials can be set using environment variables
#- AWS_DEFAULT_REGION=us-east-1
#- AWS_ACCESS_KEY_ID=<aws_access_key_id>
#- AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>
proxy:
image: delphix-hyperscale-masking-proxy:<HYPERSCALE VERSION>
init: true
networks:
  - hyperscale-net
ports:
  - "443:443"
  - "80:80"
restart: unless-stopped
depends_on:
  - controller-service
networks:
hyperscale-net:
volumes:
hyperscale-load-data:
hyperscale-unload-data:
hyperscale-masking-data:
hyperscale-controller-data

```

5. (OPTIONAL) To modify the default Hyperscale configuration properties for the application, see [Configuration Settings](#)<sup>30</sup>.

6. Run the application from the same location where you extracted the `docker-compose.yaml` file.

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


<sup>30</sup> <https://portal.document360.io/continuous-compliance/docs/configuration-settings-5>

- Run the following command to check if the application is running. The output of this command should show five containers up and running.

```
docker-compose ps
```

- Run the following command to access application logs of a given container.

```
docker logs -f service_container_name>
```


 Service container name can be accessed by output of the command `docker-compose ps`.

- Run the following command to stop the application (if required).

```
sudo docker-compose down
```

7. Once the application starts, an API key will be generated that will be required to authenticate with the Hyperscale Compliance Orchestrator. This key will be found in the docker container logs of the controller service. You can either look for the key from the controller service logs location that was set as a volume binding in the `docker-compose.yml` file or you could use the following 'docker' command to retrieve the logs.

```
docker logs -f <service_container_name>
```

 Service container name can be accessed by output of the command `docker-compose ps`.

The above command displays an output similar to the following where the string `NEWLY GENERATED API KEY` can be grepped from the log::

```
2022-05-18 12:24:10.981 INFO 7 --- [          main] o.a.c.c.C.[Tomcat].[localhost].
[/] : Initializing Spring embedded WebApplicationContext
2022-05-18 12:24:10.982 INFO 7 --- [          main]
w.s.c.ServletWebServerApplicationContext : Root WebApplicationContext: initialization
completed in 9699 ms
NEWLY GENERATED API KEY: 1.89lPH1dHSJQwHuQvzawD99sf4SpBPXJADUmJS8v00VCF4V7rjtRFAftGwy
gFfsqM
```

To authenticate with the Hyperscale Compliance Orchestrator, you must use the API key and include the HTTP Authorization request header with the type `apk`; `apk <API Key>`.

For more information, see the **Authentication** section under [Accessing the Hyperscale Compliance API \(see page 141\)](#).

### 5.5.1.2.2 Continuous Compliance Engine Installation

Delphix Continuous Compliance Engine is a multi-user, browser-based web application that provides complete, secure, and scalable software for your sensitive data discovery, masking, and tokenization needs while meeting enterprise-class infrastructure requirements. For information about installing the Continuous Compliance Engine, see [Continuous Compliance Engine Installation](#)<sup>31</sup> documentation.

### 5.5.1.3 Custom configuration

Docker Compose should only be used to deploy Hyperscale Compliance in an evaluation/testing capacity.

#### 5.5.1.3.1 Introduction

This topic provides background information on performing custom configurations that are referenced throughout Hyperscale Compliance documentation.

#### 5.5.1.3.2 Bind mounts

Configuration of Hyperscale Compliance is achieved through a combination of [Configuration Settings](#) (see [page 182](#)) and the use of Docker [bind mounts](#)<sup>32</sup>. 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.

The Hyperscale Compliance and proxy containers can both be configured via separate bind-mounted directories. Each container requires all configuration files to be mounted to the relevant 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 the relevant directory. This is done by editing the `docker-compose.yml`.

Like, 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 Hyperscale Compliance 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.

---

<sup>31</sup> <https://portal.document360.io/continuous-compliance/docs/first-time-setup>

<sup>32</sup> <https://docs.docker.com/storage/bind-mounts/>

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-hyperscale-masking-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 Hyperscale Compliance being installed).

```
docker ps
```

CONTAINER ID	IMAGE	COMMAND	CREATED
802606779b9d	delphix-hyperscale-masking-proxy:dev	"/sbin/tini -- /boot..."	3 hours ago
	Up 3 hours	0.0.0.0:80->80/tcp, :::80->80/tcp, 0.0.0.0:443->443/tcp, :::443->443/tcp	
	hyperscale-masking_proxy_1		

In the above example, `802606779b9d` 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.

## 5.5.1.4 Upgrading the Hyperscale Compliance orchestrator (Docker compose)

### 5.5.1.4.1 Pre-requisite

Before upgrading, ensure you have downloaded the Hyperscale Compliance 17.0.0 tar bundle from the Delphix [Download](https://download.delphix.com/folder/2224/Delphix%20Product%20Releases/Hyperscale%20Compliance)<sup>33</sup> website.

<sup>33</sup> <https://download.delphix.com/folder/2224/Delphix%20Product%20Releases/Hyperscale%20Compliance>

### 5.5.1.4.2 How to upgrade the Hyperscale Compliance Orchestrator

Perform the following steps to upgrade the Hyperscale Compliance Orchestrator to the 17.0.0 version:

1. Run `cd /<hyperscale_installation_path>/` and `docker-compose down` to stop and remove all the running containers.
2. Run the below commands to delete all existing dangling images and hyperscale images:

```
docker rmi $(docker images -f "dangling=true" -q)
docker rmi $(docker images "delphix-hyperscale-masking-proxy" -q)
docker rmi $(docker images "delphix-controller-service-app" -q)
docker rmi $(docker images "delphix-masking-service-app" -q)
docker rmi $(docker images "delphix-*load-service-app" -q)
```

3. Remove all files or folders from existing installation directories, except `docker-compose.yaml` (Keep its backup outside the installation directory so it is not overridden while executing the next step).
4. Take backup of `.env` file and untar the patch tar in your existing installation path. `tar -xzf delphix-hyperscale-masking-17.0.0.tar.gz -C <existing_installation_path>`
5. Replace the `docker-compose.yaml` supplied with the bundle file as per the following:
  - **For users upgrading from 3.0.0.x:** Use the connector-specific `docker-compose-sample.yaml` file (e.g. `docker-compose-oracle.yaml` or `docker-compose-mssql.yaml`) supplied with the bundle and add the same 'volumes' and/or any other properties (if configured) for each container referencing the backed-up `docker-compose.yaml` from step 3.
  - **For users upgrading from 4.0.0.0 and above:** Replace the `docker-compose.yaml` file supplied with the bundle with the `docker-compose.yaml` file that you created as a backup at step 3.
  - Similar to other services, make sure to add the volume binding for the staging area path under controller-service as well. For example,

```
- /mnt/hyperscale:/etc/hyperscale
```

6. Apply the backed up `.env` file and set the `VERSION` property as 17.0.0 (i.e. `VERSION=17.0.0`).
7. Run the below commands to load the images (will configure Oracle-based unload/load setup):

```
docker load --input controller-service.tar
docker load --input unload-service.tar
```



```
docker load --input masking-service.tar
docker load --input load-service.tar
docker load --input proxy.tar
```

- If upgrading from an MSSQL connector setup(supported starting 5.0.0.0 release), instead of running the above commands for load/unload services setup(which are for Oracle), run the below commands(rest remains same for the controller, masking, and proxy services):

```
docker load --input mssql-unload-service.tar
docker load --input mssql-load-service.tar
```

- If upgrading from a Delimited Files connector setup (supported starting 12.0.0 release), instead of running the above commands for load/unload services setup(which are for Oracle), run the below commands(rest remains same for the controller, masking, and proxy services):

```
docker load --input delimited-unload-service.tar
docker load --input delimited-load-service.tar
```


- If upgrading from a MongoDB connector setup (supported starting 13.0.0 release), instead of running the above commands for load/unload services setup (which are for Oracle), run the below commands (rest remains same for the controller, masking, and proxy services):

```
docker load --input mongodb-unload-service.tar
docker load --input mongodb-load-service.tar
```

- If upgrading from a Parquet connector setup (supported starting 17.0.0 release), instead of running the above commands for load/unload services setup (which are for Oracle), run the below commands (rest remains same for the controller, masking, and proxy services):

```
docker load --input parquet-unload-service.tar
docker load --input parquet-load-service.tar
```

8. Run `docker-compose up -d` to create containers.
9. Ensure all your mount(s) are configured and accessible, before running a job.

 Existing data remains intact after the degradation.

### 5.5.1.5 How to generate a support bundle (Docker compose)



If it is not already installed, you must install bash shell and yq to generate a Hyperscale support bundle. For more information on downloading yq, refer to the [yq Downloads page on GitHub](#)<sup>34</sup>.

#### 5.5.1.5.1 1. Find the “generate\_support\_bundle.sh” script

- Login to Hyperscale VM for which you want to generate the support bundle.
- generate\_support\_bundle.sh” file is bundled with the release tar file. You can find this script under `tools/support-scripts` folder, (present under the directory, where you will untar the release tar file on Hyperscale Engine). For example, `/path_to_untarred_hyperscale_product/tools/support-scripts`.

##### Example:

```
dlpxuser@delphix:~/test$ cd tools/support-scripts/
dlpxuser@delphix:~/test$ ls -ltr
total 48
-rwxr-xr-x  1 delphix  staff    823 Jul  7 09:55 generate_support_bundle.sh
-rwxr-xr-x  1 delphix  staff    463 Jul  7 09:55 container_information.sh
-rwxr-xr-x  1 delphix  staff   5597 Jul  7 09:55 collect_container_support_info.sh
-rw-r--r--  1 delphix  staff   5316 Jul  7 09:55 README.md
```

#### 5.5.1.5.2 2. Modify the “container\_information.sh” script parameters

Change the `mount_path` and `docker_compose_file_path` accordingly.

##### Example:

```
mount_path=/home/delphix/hyperscale
docker_compose_file_path=/home/delphix/docker-compose.yaml
```



- `mount_path`: Absolute path configured for mount directory in `docker-compose` file which is mapped to `/etc/hyperscale`.
- `docker_compose_file_path`: Absolute path for `docker_compose.yaml` file.

<sup>34</sup> [https://github.com/mikefarah/yq/#wf\\_ssl\\_install](https://github.com/mikefarah/yq/#wf_ssl_install)

### 5.5.1.5.3 3. Execute the “generate\_support\_bundle.sh” script

- Execute the “generate\_support\_bundle.sh” script from `tools/support-scripts/` folder.

**Example:**

```
dlpxuser@delphix:~/test/tools/support-scripts/$ ./generate_support_bundle.sh
....
Generating support bundle tar file...
....
```

- Enter the “Password” when prompted.

### 5.5.1.5.4 4. Find the Generated Support Bundle Tar File

The resulting support bundle will be located at `/etc/hyperscale/hyperscale-support-****.tar.gz` inside the container. This means the tar file is generated under the path which is mapped to `/etc/hyperscale` in `docker-compose` file and is directly accessible from Hyperscale VM.

**Example:**

```
dlpxuser@delphix:~/test$ ls -ltr ../hyperscale/
total 316
drwxrwxrwx 5 1004 1005 4096 Feb  9 10:14 aks-mount
-rw-r--r-- 1 65436 staff 104189 Feb 17 08:52 hyperscale-support-
<current_timestamp>.tar.gz
```

**The support bundle tar file contains the following information:**

- Hyperscale Logs
- The output of `mpstat` for CPU utilization info.
- The output of `proc/meminfo` for memory info.
- The output of `proc/cpuinfo` for cpu info.
- Files to show the memory limit for the application container and the max usage of the app container in bytes.
- Redacted database file to restore the Hyperscale VM
- Docker compose file



- The script `generate_support_bundle.sh` generates a bare-bones support bundle from a Hyperscale engine running in docker.
- Execute the `generate_support_bundle.sh` from the untar location.

- The resulting support bundle will be at `/etc/hyperscale/hyperscale-support-****.tar.gz` inside the container. This means the tar file is generated under a path that is mapped to `/etc/hyperscale` in `docker-compose` file and is directly accessible from Hyperscale VM.
- The user should have privileges or permission to execute the docker command in order to generate the support bundle.

### 5.5.1.6 Managing the storage space

There are two storage locations where the continuous increase in disk space consumption can lead to a depletion of the available disk space on the Hyperscale host system.

#### 5.5.1.6.1 1. Overlay2 File System

The `/var/lib/docker/overlay2` directory stores several [file system layers](#)<sup>35</sup> for images and containers, and it may also accumulate data related to unused containers and images.

To prevent excessive disk space usage caused by the Overlay2 File System, it's essential to follow the Hyperscale upgrade steps, including the deletion of dangling images as specified. If this storage becomes full, run the below command:

```
docker rmi $(docker images -f "dangling=true" -q)
```

#### 5.5.1.6.2 2. Container Logs

Docker maintains logs printed on the console by each service in JSON format within the `/var/lib/docker/containers` directory (because the default driver is json-file). These logs keep appending to these files. If this storage becomes full, perform the below steps:

1. Create a `daemon.json` file under `/etc/docker` directory (For more information, refer to <https://docs.docker.com/config/containers/logging/local/#usage%5D> & [Configure logging driver](#)<sup>36</sup>). The below example describes the sample content of the file.

```
{
  "log-driver": "json-file",
  "log-opts": {
    "max-size": "2g",
    "max-file": "3"
  }
}
```

<sup>35</sup> <https://docs.docker.com/storage/storagedriver/>

<sup>36</sup> <https://docs.docker.com/config/containers/logging/configure/#configure-the-default-logging-driver>

2. Reload & restart the docker service.

```
sudo systemctl daemon-reload
sudo systemctl restart docker
```

3. Stop all currently running docker services & then restart these again. Execute the below commands from the HS directory.

```
docker-compose down
docker-compose up -d
```

### 5.5.1.7 Migrating to Kubernetes

#### 1. Copy docker-compose services data to local storage

Stop Hyperscale services so that all of the product's state is flushed to persistent storage. You must run the below command from the same directory where the `docker-compose.yaml` file exists.

```
docker-compose stop
```

Create folders for data backup of all services.

```
mkdir container_data
cd container_data/
mkdir controller
mkdir unload
mkdir masking
mkdir load
mkdir proxy
cd ../
```

Copy the SQLite database file and encryption key Docker volume folder data on a local machine in the created folder.

```
docker ps -a // list out the name of containers

docker cp <controller-service-container-name>:/data ./container_data/controller/
docker cp <unload-service-container-name>:/data ./container_data/unload/
docker cp <load-service-container-name>:/data ./container_data/load/
docker cp <masking-service-container-name>:/data ./container_data/masking/

eg: docker cp hyperscale-masking_controller-service_1:/data ./container_data/
controller/
```

Copy the SSL certificates Docker volume folder data on a local machine in the created folder (This step is only required if custom certificates are used).

```
docker cp <proxy-container-name>:/etc/config/nginx/ssl ./container_data/proxy
```

All the `container_data` files should be as below.

```

├── controller
│   └── data
│       └── db.sqlite
├── load
│   └── data
│       ├── db.sqlite
│       └── encryption.key
├── masking
│   └── data
│       ├── db.sqlite
│       └── encryption.key
├── proxy
│   └── ssl
│       ├── dhparam.pem
│       ├── nginx.crt
│       ├── nginx.key
│       └── ssl.conf
└── unload
    └── data
        ├── db.sqlite
        └── encryption.key

```

Grant permission 770 to all files.

```
chmod 770 -R .
```



Please note that docker-compose log files will not be migrated so keep the backup, of log files before starting the migration process.

## 2. Restore docker-compose services data to Hyperscale deployed on the Kubernetes setup (in Persistent Volume)

Install the required version of the Hyperscale Compliance Engine using Kubernetes. For more information about the installation steps, refer to the <https://hyperscalemasking.delphix.com/docs/latest/installation-and-setup-kubernetes> section.

List the pods name for further reference:

```
kubectl get pods -n hyperscale-services
```

Restore Hyperscale docker-compose version volume data with Hyperscale deployed on the Kubernetes setup (in Persistent Volume).

Execute the below process for all the pods names except the proxy pod to copy each service data.

```
cd container_data/<service- folder-name>
kubectl cp data hyperscale-services/<service-pod-name>:/

eg:
cd container_data/masking
kubectl cp data hyperscale-services/masking-service-7788ccbbb-lnzhv:/
```

### 3. List out deployment names and restart all the services using kubectl rollout.

```
kubectl get deployments -n hyperscale-services
kubectl rollout restart deployment <controller-service-deployment-name> -n
hyperscale-services
eg: kubectl rollout restart deployment controller-service -n hyperscale-services
```

### 4. Update the SSL certificates (This step is only required if custom certificates are used).

Go to `container_data/proxy/ssl` folder and get the base64 value of `nginx.crt`, `nginx.key`, and `dhparam.pem`. For example:

```
cat nginx.crt | base64 | awk '{print}' ORS='' | awk '{print}'
```

Now update the base64 value of each in k8 `values.yaml` file.

### 5. Update the other user-configurable properties from docker-compose .env file to values.yaml accordingly

All the available configurations can be found at <https://hyperscalemasking.delphix.com/docs/latest/configuration-settings>, the configuration migration is divided into two steps:

#### 5.5.1.7.1 Modify existing configuration for any service:

Go to the `hyperscale-helm` chart directory. Under the templates folder, locate the corresponding `<service-name>-deployment.yaml` file. For more details, refer to the below screenshot.

Name	Date Modified	Size	Kind
Chart.yaml	01-Dec-2023, 7:33 PM	129 bytes	YAML Document
README.md	01-Dec-2023, 7:33 PM	39 bytes	Markdo...ument
templates	02-Dec-2023, 1:34 PM	--	Folder
controller-database-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	330 bytes	YAML Document
<b>controller-deployment.yaml</b>	<b>01-Dec-2023, 7:33 PM</b>	<b>4 KB</b>	<b>YAML Document</b>
controller-service.yaml	01-Dec-2023, 7:33 PM	317 bytes	YAML Document
docker-registry-secret.yaml	01-Dec-2023, 7:33 PM	258 bytes	YAML Document
hyperscale-net-networkpolicy.yaml	01-Dec-2023, 7:33 PM	325 bytes	YAML Document
hyperscale-services-namespace.yaml	01-Dec-2023, 7:33 PM	176 bytes	YAML Document
imagePullSecret.tpl	01-Dec-2023, 7:33 PM	352 bytes	Document
instantclient-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	530 bytes	YAML Document
instantclient-volume.yaml	01-Dec-2023, 7:33 PM	640 bytes	YAML Document
load-database-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	304 bytes	YAML Document
<b>load-deployment.yaml</b>	<b>01-Dec-2023, 7:33 PM</b>	<b>3 KB</b>	<b>YAML Document</b>
load-service.yaml	01-Dec-2023, 7:33 PM	413 bytes	YAML Document
masking-database-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	313 bytes	YAML Document
<b>masking-deployment.yaml</b>	<b>01-Dec-2023, 7:33 PM</b>	<b>2 KB</b>	<b>YAML Document</b>
masking-service.yaml	01-Dec-2023, 7:33 PM	422 bytes	YAML Document
proxy-certificate.yaml	01-Dec-2023, 7:33 PM	363 bytes	YAML Document
<b>proxy-deployment.yaml</b>	<b>01-Dec-2023, 7:33 PM</b>	<b>2 KB</b>	<b>YAML Document</b>
proxy-service.yaml	01-Dec-2023, 7:33 PM	276 bytes	YAML Document
stage-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	604 bytes	YAML Document
stage-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	440 bytes	YAML Document
unload-database-persistentvolumeclaim.yaml	01-Dec-2023, 7:33 PM	336 bytes	YAML Document
<b>unload-deployment.yaml</b>	<b>01-Dec-2023, 7:33 PM</b>	<b>3 KB</b>	<b>YAML Document</b>
unload-service.yaml	01-Dec-2023, 7:33 PM	419 bytes	YAML Document
tools	02-Dec-2023, 1:34 PM	--	Folder
values.yaml	01-Dec-2023, 7:33 PM	8 KB	YAML Document

Open the corresponding `<service-name>-deployment.yaml` file for which you want to modify the new configuration.

**Example:** Let's suppose you want to change the config value for `API_KEY_CREATE` in `controller-deployment.yaml` file. As you can see that `API_KEY_CREATE` config value is mapped with `Values.controller.apiKeyCreate`.



```

1 #
2 # Copyright (c) 2023 by Delphix. All rights reserved.
3 #
4
5 apiVersion: apps/v1
6 kind: Deployment
7 metadata:
8   labels:
9     io.hyperscale.service: controller-service
10   name: controller-service
11   namespace: {{ .Values.namespace }}
12 spec:
13   replicas: 1
14   selector:
15     matchLabels:
16       io.hyperscale.service: controller-service
17   strategy: {}
18   template:
19     metadata:
20       labels:
21         io.hyperscale.service: controller-service
22     spec:
23       {{- if .Values.nodeName }}
24       nodeName: {{ .Values.nodeName }}
25       {{- end }}
26     securityContext:
27       runAsUser: 65436
28       runAsGroup: 50
29       fsGroup: 50
30     containers:
31     - env:
32       - name: API_KEY_CREATE
33         value: {{ .Values.controller.apiKeyCreate | quote }}
34       - name: API_VERSION_COMPATIBILITY_STRICT_CHECK
35         value: {{ .Values.controller.apiVersionCompatibilityStrictCheck | quote }}
36       - name: EXECUTION_STATUS_POLL_DURATION
37         value: {{ .Values.controller.executionStatusPollDuration | quote }}
38       - name: LOAD_SERVICE_REQUIREPOSTLOAD
39         value: {{ .Values.controller.loadServiceRequirepostload | quote }}
40       - name: LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE
41         value: {{ .Values.controller.loggingLevelRoot | quote }}
42       - name: SKIP_UNLOAD_SPLIT_COUNT_VALIDATION
43         value: {{ .Values.controller.skipUnloadSplitCountValidation | quote }}
44       - name: SKIP_LOAD_SPLIT_COUNT_VALIDATION
45         value: {{ .Values.controller.skipLoadSplitCountValidation | quote }}
46       - name: MASKING_SERVICE_BASE_URL
47         value: {{ .Values.maskingServiceBaseUrl }}
48       - name: UNLOAD_SERVICE_BASE_URL
49         value: {{ .Values.unloadServiceBaseUrl }}
50       - name: LOAD_SERVICE_BASE_URL
51         value: {{ .Values.loadServiceBaseUrl }}
52       - name: CANCEL_STATUS_POLL_DURATION
53         value: {{ .Values.controller.cancelStatusPollDuration | quote }}
54       # uncomment below lines for MongoDB / Delimited connector
55       # name: SOURCE_KEY_FIELD_NAMES
56       # value: {{ .Values.controller.sourceKeyFieldNames | quote }}
57       # uncomment below lines only for MongoDB connector
58       # name: VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS
59       # value: {{ .Values.controller.validateUnloadRowCountForStatus | quote }}
60       # name: VALIDATE_MASKED_ROW_COUNT_FOR_STATUS
61       # value: {{ .Values.controller.validateMaskedRowCountForStatus | quote }}
62       # name: VALIDATE_LOAD_ROW_COUNT_FOR_STATUS
63       # value: {{ .Values.controller.validateLoadRowCountForStatus | quote }}
64       # name: DISPLAY_BYTES_INFO_DL_STATUS
65       # value: {{ .Values.controller.displayBytesInfoInStatus | quote }}
66       # name: DISPLAY_ROW_COUNT_IN_STATUS
67       # value: {{ .Values.controller.displayRowCountInStatus | quote }}
68

```

Now under the `hyperscale-helm` chart directory, open the `values.yaml` file and search for `apiKeyCreate` property of controller attribute and change the corresponding value.

### 5.5.1.7.2 Add a new configuration for any service

In case the configuration name doesn't exist in `<service-name>-deployment.yaml` file, add your configuration as part of `env` in `<service-name>-deployment.yaml` file and also define config with the same name in `values.yaml` corresponding to the service environment as explained above, to get the values from `values.yaml` file.

Apply helm upgrade to reflect the new `values.yaml` file changes.

```
helm upgrade hyperscale-helm -f hyperscale-helm/values.yaml hyperscale-helm
```

Your docker-compose Hyperscale VM is now fully migrated to Kubernetes Hyperscale VM. Check the k8 running pods, like below:

```
delphix@hyperscale-k8s:~/_database$ kubectl get pods -n hyperscale-services
NAME                                READY   STATUS    RESTARTS   AGE
masking-service-7788ccbbb-lnzhv    1/1     Running   0           4h28m
unload-service-5f9946556f-4qg7l    1/1     Running   0           4h28m
proxy-767d6ccbdd-2vdhv             1/1     Running   0           4h28m
load-service-5bb994dbbf-c85mz      1/1     Running   0           3h43m
controller-service-6c6df56bbd-blng6 1/1     Running   0           62m
```

If `apiKeyCreate` property is defined as true in `values.yaml`, then you can get the API-Key as below to access the API. In case, `apiKeyCreate` property is defined as false then you can use your existing API-Key.

```
kubectl logs <controller-pod-name> -n hyperscale-services | grep 'NEWLY GENERATED API KEY' | sed 's/^.*: //'
```

## 5.5.2 Kubernetes

This section covers the following topics:

- [Host requirements \(Kubernetes\)](#) (see page 106)
- [Docker image registry and names for Hyperscale services](#) (see page 107)
- [Installation and setup \(Kubernetes\)](#) (see page 109)
- [Bootstrapping API keys \(Kubernetes\)](#) (see page 115)
- [Hyperscale logs \(Kubernetes\)](#) (see page 116)
- [Limitations \(Kubernetes\)](#) (see page 116)
- [Upgrading the Hyperscale Compliance orchestrator \(Kubernetes\)](#) (see page 116)
- [How to generate a support bundle \(Kubernetes\)](#) (see page 118)

### 5.5.2.1 Host requirements (Kubernetes)

Type	Host Requirement	Explanation
User	A user (example: <code>hyperscale_os</code> ) with the following permissions is required: <ul style="list-style-type: none"> <li>• Permission to run helm commands</li> <li>• Permission to run <code>kubectl</code> commands</li> </ul>	These permissions are required to be able to install helm charts and manage the Kubernetes resources.
NFS Client Services	NFS client services must be enabled on the host.	NFS client service is required to be able to mount an NFS shared storage from where the Hyperscale Compliance Orchestrator will be able to read the

Type	Host Requirement	Explanation
		source files and write the target files. For more information, see <a href="#">NFS Server Installation</a> <sup>37</sup> .
Staging area directory permissions	<p>The 'staging_area' directory, along with the directory( <code>hyperscale</code> ) one level above, require the following permissions, based on the UID/GID of the OS user so that the Hyperscale Compliance Orchestrator and the Continuous Compliance Engine(s) can perform read/write/execute operations on the staging area:</p> <ul style="list-style-type: none"> <li>• If the Hyperscale Compliance OS user has a UID of 65436, then the 'staging_area' directory, along with the directory( <code>hyperscale</code> ) one level above, must have a UID of 65436 and 700 permission mode.</li> <li>• If the Hyperscale Compliance OS user has a GID of 50 and does not have a UID of 65436, then the 'staging_area' directory, along with the directory( <code>hyperscale</code> ) one level above, must have a GID of 50 and 770 permission mode.</li> </ul>	These permissions on the staging area directory are required for the Hyperscale Compliance containers and the continuous compliance engines to be able to read/write into the staging area.
Installation Directory	A directory to download and manage helm charts.	
Hardware Requirements	<p><b>Minimum:</b> 8 vCPU, 64 GB of memory, 50GB OS disk.</p> <p><b>Recommended:</b> 16 vCPU, 128GB of memory, 50GB OS disk.</p>	

### 5.5.2.2 Docker image registry and names for Hyperscale services

The docker images for Hyperscale Compliance services fall under two categories of services:

1. **Common Services:** A set of service images common across deployments with varying unload/load services.
2. **Unload and Load Services:** A set of unload and load service images specific to a dataset vendor.

<sup>37</sup> <http://delphixdocs.atlassian.net/hyperscale-compliance-9-0-0/docs/nfs-server-installation-3>

All images are stored under the [hyperscale.download.delphix.com/delphix-hyperscale](https://hyperscale.download.delphix.com/delphix-hyperscale)<sup>38</sup> registry. Also, any image will have a name of the form `<image-name>-[VERSION]` where [VERSION] corresponds to a particular release version.

The following list provides the image URLs for the docker images in the two categories of services:

#### 5.5.2.2.1 Common Services

```
hyperscale.download.delphix.com/delphix-hyperscale:proxy-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:controller-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:masking-service-[VERSION]
```

#### 5.5.2.2.2 Unload and Load Services

##### 5.5.2.2.2.1 Oracle Unload and Load

```
hyperscale.download.delphix.com/delphix-hyperscale:oracle-unload-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:oracle-load-service-[VERSION]
```

##### 5.5.2.2.2.2 MSSQL Unload and Load:

```
hyperscale.download.delphix.com/delphix-hyperscale:mssql-unload-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:mssql-load-service-[VERSION]
```

##### 5.5.2.2.2.3 Delimited Unload and Load:

```
hyperscale.download.delphix.com/delphix-hyperscale:delimited-unload-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:delimited-load-service-[VERSION]
```

##### 5.5.2.2.2.4 Mongo Unload and Load:

```
hyperscale.download.delphix.com/delphix-hyperscale:mongo-unload-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:mongo-load-service-[VERSION]
```

---

<sup>38</sup> <http://hyperscale.download.delphix.com/delphix-hyperscale>

#### 5.5.2.2.5 Parquet Unload and Load:


```
hyperscale.download.delphix.com/delphix-hyperscale:parquet-unload-service-[VERSION]
hyperscale.download.delphix.com/delphix-hyperscale:parquet-load-service-[VERSION]
```

### 5.5.2.3 Installation and setup (Kubernetes)

#### 5.5.2.3.1 Installation Requirements

To deploy Hyperscale Compliance via Kubernetes, a running Kubernetes cluster is required to run, the `kubectl` command line tool to interact with the Kubernetes cluster and HELM for deployment onto the cluster.

Requirement	Recommended Version	Comments
Kubernetes Cluster	1.25 or above	
HELM	3.9.0 or above	HELM installation should support HELM v3. More information on HELM can be found at <a href="https://helm.sh/docs/">https://helm.sh/docs/</a> . To install HELM, follow the installation instructions at <a href="https://helm.sh/docs/intro/install/">https://helm.sh/docs/intro/install/</a> .  The installation also requires access to the HELM repository from where Hyperscale charts can be downloaded. The HELM repository URL is <a href="https://dlpx-helm-hyperscale.s3.amazonaws.com">https://dlpx-helm-hyperscale.s3.amazonaws.com</a> .
kubectl	1.25.0 or above	

 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 the **imageCredentials** section.

### 5.5.2.3.2 Installation

#### 5.5.2.3.2.1 Download the HELM charts

The latest version of the chart can be pulled locally with the following command:

```
curl -XGET https://dlpx-helm-hyperscale.s3.amazonaws.com/hyperscale-helm-17.0.0.tgz -o hyperscale-helm-17.0.0.tgz
```

This command will download a file with the name `hyperscale-helm-17.0.0.tgz` in the current working directory. The downloaded file can be extracted using the following command:

```
tar -xvf hyperscale-helm-17.0.0.tgz
```

This will extract into the following directory structure:

```
hyperscale-helm
|- values.yaml
|- README.md
|- Chart.yaml
|- templates
  |-<all templates files>
```

#### 5.5.2.3.2.2 Verify the authenticity of the downloaded HELM charts

The SHA-256 hash sum of the downloaded helm chart tarball file can be verified as follows:

1. Execute the below command and note the digest value for version 17.0.0  

```
curl https://dlpx-helm-hyperscale.s3.amazonaws.com/index.yaml
```
2. Execute the `sha256sum` command (or equivalent) on the downloaded file(`hyperscale-helm-17.0.0.tgz`)  

```
sha256sum hyperscale-helm-17.0.0.tgz
```


The value generated by the `sha256sum` utility in step 2 must match the digest value noted in step 1.

#### 5.5.2.3.2.3 Configure Registry Credentials for Docker Images

For pulling the Docker images from the registry, permanent credentials associated with your Delphix account would need to be configured in the `values.yaml` file. To get these permanent credentials, visit the Hyperscale Compliance Download page and log in with your credentials. Once logged in, select the Hyperscale HELM Repository link and accept the Terms and Conditions. Once accepted, credentials for the docker image registry will be presented. Note them down and edit the `imageCredentials.username` and `imageCredentials.password` properties in the `values.yaml` file as shown below:

```
# Credentials to fetch Docker images from Delphix internal repository
  imageCredentials:
# Username to login to docker registry
  username: <username>
# Password to login to docker registry
  password: <password>

imageCredentials:
username: <username>
password: <password>
```

 Delphix will delete unused credentials after 30 days and inactive (but previously used) credentials after 90 days.

#### 5.5.2.3.2.4 Override Default Values in values.yaml

`hyperscale-helm` is the name of the folder which was extracted in the previous step. In the above directory structure, the `values.yaml` file contains all of the configurable properties with their default values. These default values can be overridden while deploying Hyperscale Compliance, as per the requirements. If the `values.yaml` file needs to be overridden, create a copy of `values.yaml` and edit the required properties. While deploying Hyperscale Compliance, `values.yaml` file can be overridden using either of the following commands:

```
helm install hyperscale-helm -f <path to edited values.yaml> <directory path of the
extracted chart>
helm install hyperscale-helm <directory path of the extracted chart> --set
<property1>=<value1> --set <property2>=<value2>
```

#### 5.5.2.3.2.5 Configure Helm chart properties of importance

A few commonly used properties for each service (controller, masking, unload and load) with their default values are listed in `values.yaml`. You may find details of these properties on the [Configuration Settings](#)<sup>39</sup> page. The following sections talk about some of the important properties that will need to be configured correctly for a successful deployment.

<sup>39</sup> <http://delphixdocs.atlassian.net/hyperscale-compliance-9-0-0/docs/configuration-settings-7>

### 5.5.2.3.2.6 Configure the staging area

By default, a path( `/d\pxdata` ) local to the Kubernetes cluster node will be used, via [persistent volume](#)<sup>40</sup> claims, to mount the staging area path inside the pods. Override the path by setting up the desired local storage path with the **`localStoragePath`** property.

If the cluster needs to mount an NFS shared path that will act as the staging area, override the **`nfsStorageHost`** and **`nfsStorageExportPath`** properties.

### 5.5.2.3.2.7 Configure the correct unload/load images for your dataset type

By default, the helm installation will create the Kubernetes pods for the Oracle unload and load services. To have helm create unload and load pods for the other connector types, for example, MSSQL unload and load services, override the following values in the `values.yaml` file:

1. `unload.imageName=mssql-unload-service`
2. `load.imageName=mssql-load-service`

For Delimited File connector, change the following in `values.yaml` file:

1. `unload.imageName=delimited-unload-service`
2. `load.imageName=delimited-load-service`

For MongoDB Database connector, change the following in `values.yaml`

1. `unload.imageName=mongo-unload-service`
2. `load.imageName=mongo-load-service`

For the Parquet connector, change the following in `values.yaml` file:

1. `unload.imageName=parquet-unload-service`
2. `load.imageName=parquet-load-service`

For more information about the docker image names and their URLs, refer to [Docker image registry and names for Hyperscale services](#) (see page 107).

### 5.5.2.3.2.8 Configure controller-service environment variables for the MongoDB Database Connector (Applicable only for the MongoDB Database masking)

- Change the following in `values.yaml` :
- controller:
  - `controller.loadServiceRequirepostload=false`

---

<sup>40</sup> <https://kubernetes.io/docs/concepts/storage/persistent-volumes/>



- `controller.SourceKeyFieldNames=database_name,collection_name`
- `controller.validateUnloadRowCountForStatus: false`
- `controller.validateMaskedRowCountForStatus: false`
- `controller.validateLoadRowCountForStatus: false`
- `controller.displayBytesInfoInStatus: true`
- `controller.displayRowCountInStatus: false`
- unload:
  - `unload.concurrentExportLimit: 10`
- Navigate to `hyperscale-helm/templates/controller-deployment.yaml` and uncomment lines:

```
spec:
  ...
  template:
    ...
    containers:
      - env:
          ...
          - name: SOURCE_KEY_FIELD_NAMES
            value: {{ .Values.controller.sourceKeyFieldNames | quote }}
          - name: VALIDATE_UNLOAD_ROW_COUNT_FOR_STATUS
            value: {{ .Values.controller.validateUnloadRowCountForStatus |
quote }}
          - name: VALIDATE_MASKED_ROW_COUNT_FOR_STATUS
            value: {{ .Values.controller.validateMaskedRowCountForStatus |
quote }}
          - name: VALIDATE_LOAD_ROW_COUNT_FOR_STATUS
            value: {{ .Values.controller.validateLoadRowCountForStatus |
quote }}
          - name: DISPLAY_BYTES_INFO_IN_STATUS
            value: {{ .Values.controller.displayBytesInfoInStatus | quote }}
          - name: DISPLAY_ROW_COUNT_IN_STATUS
            value: {{ .Values.controller.displayRowCountInStatus | quote }}
```

- Navigate to `hyperscale-helm/templates/unload-deployment.yaml` and uncomment lines:


```
spec:
  ...
  template:
    ...
    containers:
      - env:
          ...
          - name: CONCURRENT_EXPORT_LIMIT
            value: {{ .Values.unload.concurrentExportLimit | quote }}
```

### 5.5.2.3.2.9 Configure the instantclient path(Applicable for Oracle unload/load):

By default, a path( `/d\pxdata` ) local to the Kubernetes cluster node will be used, via [persistent volume](#)<sup>41</sup> claims, to mount the Oracle **instantclient** path inside the pods. Override the path by setting up the desired **instantclient** path with the **instantClientStoragePath** and **instantClientRootDirName** properties.

If the cluster needs to mount an NFS shared path that will contain the **instantclient** binaries, override the **nfsInstantClientHost** and **nfsInstantClientExportPath** properties.

Service level properties:

-  • HELM will internally refer to the kubeconfig file to connect to the Kubernetes cluster. The default kubeconfig file is present at location: `~/.kube/config`.
- If the kubeconfig file needs to be overridden while running HELM commands, set the **KUBECONFIG** environment variable to the location of the kubeconfig file.
- Oracle Load doesn't support Object Identifiers(OIDs).

### 5.5.2.3.2.10 Configuring the Ingress Controller

Assuming an ingress controller configuration on the Kubernetes cluster is present when accessing Hyperscale Compliance 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).

Additionally, the following annotations will need to be set(this assumes that Kubernetes [Ingress NGINX Controller](#)<sup>42</sup> is being used):

1. `nginx.ingress.kubernetes.io/backend-protocol=HTTPS`
2. `nginx.ingress.kubernetes.io/proxy-body-size=50m`
3. `nginx.ingress.kubernetes.io/proxy-connect-timeout=600`
4. `nginx.ingress.kubernetes.io/proxy-read-timeout=600`
5. `nginx.ingress.kubernetes.io/proxy-send-timeout=600`

If an ingress controller has not been assigned, then a new ingress resource, with the above requirements, can be created with the following `kubectl` command:

```
kubectl create ingress https-ingress --namespace=<namespace-name> --rule="/
*=>proxy:443" --annotation=nginx.ingress.kubernetes.io/backend-protocol=HTTPS --
annotation=nginx.ingress.kubernetes.io/proxy-body-size=50m --
annotation=nginx.ingress.kubernetes.io/proxy-connect-timeout=600 --
```

<sup>41</sup> <https://kubernetes.io/docs/concepts/storage/persistent-volumes/>

<sup>42</sup> <https://github.com/kubernetes/ingress-nginx>

```
annotation=nginx.ingress.kubernetes.io/proxy-read-timeout=600 --
annotation=nginx.ingress.kubernetes.io/proxy-send-timeout=600
```

### 5.5.2.3.2.11 Check for the Successful Installation

After installing the helm chart and setting up the ingress controller, check the status of the helm chart and the pods using the following commands:

```
$ helm list
NAME                NAMESPACE    REVISION    UPDATED
STATUS             CHART        APP VERSION
hyperscale-helm     default      1           2023-04-17 05:38:17.639357049 +0000 UTC
deployed           hyperscale-helm-17.0.0
```

```
$ kubectl get pods --namespace=hyperscale-services

NAME                                READY   STATUS    RESTARTS   AGE
controller-service-65575b6458-2q9b4 1/1     Running   0           125m
load-service-5c644b9cc8-g9fs8        1/1     Running   0           125m
masking-service-7ddfd49c8f-5j2q5     1/1     Running   0           125m
proxy-5bd8d8f589-gkx8g               1/1     Running   0           125m
unload-service-55b5bd8cc8-7z95b      1/1     Running   0           125m
```

### 5.5.2.4 Bootstrapping API keys (Kubernetes)

Once the application starts, an API key will be generated that will be required to authenticate with the Hyperscale Compliance Orchestrator. This key will be found in the logs of the controller service pod. You can use the following command to get the API key:

```
kubectl logs <controll_service_pod_name> -n <namespace> | grep 'NEWLY GENERATED
API KEY'
```

The above command displays an output similar to the following:

```
NEWLY GENERATED API KEY:
1. bTYUvuzXgnhS8U7WwYZKyF27eg01B73pJUxyw2fHAXhgVLweMIFB6L0isfA3ZGNI
```

To authenticate with the Hyperscale Compliance Orchestrator, you must use the API key and include the HTTP Authorization request header with type apk; apk <API Key> ;.

For more information, see the **Authentication** section under [Accessing the Hyperscale Compliance API](#) (see page 141).

After acquiring the bootstrap API key, you can [create \(see page 0\)](#) a new API key for your convenience. After generating the new API key, you should override the `apiKeyCreate` property in the installed helm chart by running the following command.

```
helm upgrade <release_name> <directory path of the extracted chart> --reuse-values --set apiKeyCreate=false
```

### 5.5.2.5 Hyperscale logs (Kubernetes)

All Hyperscale Compliance containers log to **`stdout`** and **`stderr`** so that their logs are processed by Kubernetes. To view container-level logs running on the Kubernetes cluster, run the following command:

```
kubectl logs <pod_name> -n <namespace>
```

### 5.5.2.6 Limitations (Kubernetes)

Hyperscale Compliance 17.0.0 does not support deploying/scaling up the services to multiple nodes in a Kubernetes cluster. This support will be added in a future release. Currently, Hyperscale Compliance is only supported on single-node Kubernetes clusters.

### 5.5.2.7 Upgrading the Hyperscale Compliance orchestrator (Kubernetes)

Perform the following steps to upgrade a Hyperscale Compliance Orchestrator (Kubernetes).

1. Create a new folder called `hyperscale-helm-[version]`, where `[version]` is the latest version to which the platform is being upgraded.

```
$ mkdir hyperscale-helm-[version]
```

2. Download the new version of the chart using the following command in tandem with the newly created folder. **Note:** This command will download a file named `hyperscale-helm-[version].tgz` in the folder `hyperscale-helm-[version]`.

```
$ cd hyperscale-helm-[version]
$ curl -XGET https://dlpx-helm-hyperscale.s3.amazonaws.com/hyperscale-helm-[version].tgz -o hyperscale-helm-[version].tgz
```

3. The downloaded file is then extracted using the following command.

```
$ tar -xvf hyperscale-helm-[version].tgz
```

4. This will extract into the following directory structure.

```
hyperscale-helm
|- values.yaml
|- README.md
```

```
|- Chart.yaml
|- templates
  |-<all templates files>
```

5. Copy the `values.yaml` file from the previous version parallel to the `hyperscale-helm-[version]` folder.
6. After copying the `values.yaml` file, there are updates that need to be made to the file under the `imageCredentials` section:
  - Bumping up the version, specified against the `tag` property, to the desired higher version.
  - If the credentials configured in your `values.yaml` have expired, which will be the case if the credentials were unused for 30 days or were inactive (but previously used) for 90 days, please retrieve a new set of credentials by visiting the [Hyperscale Compliance](http://download.delphix.com) page on <http://download.delphix.com> and selecting the [Hyperscale Helm Repository](#) link. Configure your `values.yaml` with the new set of credentials.
    - This password update in `values.yaml` is only required if the user is using a Delphix-provided Docker Registry ([hyperscale.download.delphix.com/delphix-hyperscale](http://hyperscale.download.delphix.com/delphix-hyperscale)<sup>43</sup>) directly in the deployment (i.e. `values.yaml`).
    - If you are using your internal Docker Registry, you should first pull the next version of the Docker images from the Delphix-provided registry with the credentials associated with your Delphix account.
    - The following are the 'docker' commands that can be used to pull Docker images into your internal Docker Registry:
      - Docker login command (username and password are your permanent credentials retrieved from [Delphix Download](https://download.delphix.com)<sup>44</sup> site).

```
$ docker login --username [USER] --password [PASSWORD]
```

- Pull Docker images of the Hyperscale Compliance services.

```
$ docker pull <image-url>
```

The image URLs for the Hyperscale Compliance services can be referenced from this [page](#) (see [page 107](#)) of the documentation.


7. Run the helm upgrade command.

```
$ helm upgrade -f values.yaml hyperscale-helm hyperscale-helm
```

<sup>43</sup> <http://hyperscale.download.delphix.com/delphix-hyperscale>

<sup>44</sup> <https://download.delphix.com>

### 5.5.2.8 How to generate a support bundle (Kubernetes)

 If it is not already installed, you must install bash shell and yq to generate a Hyperscale support bundle. For more information on downloading yq, refer to [yq Downloads page on GitHub](#)<sup>45</sup>.

#### 5.5.2.8.1 1. Find the “generate\_support\_bundle.sh” script

- Login to Hyperscale VM for which you want to generate the support bundle.
- generate\_support\_bundle.sh” file is bundled with the helm package file. You can find this script under the `tools/` directory (present under the directory, where you will untar the helm package tar file). For example, `/path_to_hyperscale_helm/tools/`.

**Example:**

```
dlpxuser@delphix:~/test$ cd tools/support-scripts/
dlpxuser@delphix:~/test$ ls -ltr
total 16
-rw-r--r--  1 delphix staff  1032 May  3 13:09 generate_support_bundle.sh
-rw-r--r--  1 delphix  staff    54 May  3 13:12 values-redact.properties
```

#### 5.5.2.8.2 2. Modify the “values-redact.properties”

The `values-redact.properties` file contains property values (in `values.yaml` file) which are sensitive and should be redacted before adding them to the support bundle. By default, it includes image repository-related properties. The property name will follow a format like `.<rootProperty>.<childProperty>.<innerChildProperty>`.

**Example:**

```
.unload.loggingLevelRoot // here we want to redact loggingLevelRoot property of
unload
```

#### 5.5.2.8.3 3. Execute the “generate\_support\_bundle.sh” script

Execute the “generate\_support\_bundle.sh” script from `hyperscale-helm/tools/` directory.

**Example:**

---

<sup>45</sup> [https://github.com/mikefarah/yq/#wf\\_ssl\\_install](https://github.com/mikefarah/yq/#wf_ssl_install)

```
delphix@ip-xx-xxx-xxx-xxx:~$ ./hyperscale-helm/tools/generate_support_bundle.sh
/home/delphix/support/values.yaml created.
.....
Generating support bundle tar file...
.....
```

#### 5.5.2.8.4 4. Find the Generated Support Bundle Tar File

The resulting support bundle ( `hyperscale-support-****.tar.gz` ) will be located at the same level where the `hyperscale-helm` was extracted

**Example:**

```
delphix@ip-10-110-254-92:~$ ls
hyperscale-helm  hyperscale-helm-9.0.0-6.tgz  hyperscale-support-20230502-14-36-36.tar.gz
```

**The support bundle tar file contains the following information:**

- Hyperscale Logs
- The output of `mpstat` for CPU utilization info.
- The output of `proc/meminfo` for memory info.
- The output of `proc/cpuinfo` for CPU info.
- Files to show the application container's memory limit and the app container's max usage in bytes.
- Redacted database file to restore the Hyperscale VM
- Redacted `values.yaml` file

### 5.5.3 Podman compose



Delphix highly recommends new installations be performed on Kubernetes.

This section covers the following topics:

- [Host requirements \(Podman Compose\)](#) (see page 120)
- [Installation and Setup \(Podman Compose\)](#) (see page 121)
- [How to Generate a Support Bundle \(Podman Compose\)](#) (see page 136)

### 5.5.3.1 Host requirements (Podman Compose)

Type	Host Requirement	Explanation
User	<p>A user (hyperscale_os) with the following permissions are required:</p> <ul style="list-style-type: none"> <li>• Should have permissions to install <code>podman</code> and <code>podman-compose</code>.</li> <li>• Permission to run <code>mount</code>, <code>umount</code>, <code>mkdir</code> and <code>rmdir</code> as a super-user with <code>NOPASSWD</code>.</li> <li>• sudo permission to run <code>sudo sysctl net.ipv4.ip_unprivileged_port_start=80</code></li> <li>• Should have either <code>GID=50</code> and/or <code>UID=65436</code>.</li> </ul>	<p>This will be a primary user responsible to install and operate the Hyperscale Compliance.</p>
Installation Directory	<p>There must be a directory on the Hyperscale Compliance Orchestrator host where the Hyperscale Compliance can be installed.</p>	<p>This is a directory where the Hyperscale Compliance tar archive file will be placed and extracted. The extracted artifacts will include docker images(tar archive files) and a configuration file(<code>podman-compose.yaml</code>) that will be used to install the Hyperscale Compliance.</p>
Log File Directory	<p>An optional directory to place log files.</p>	<p>This directory (can be configured via <code>podman-compose.yaml</code> configuration file) will host the runtime/log files of the Hyperscale Compliance Orchestrator.</p>
NFS Client Services	<p>NFS client services must be enabled on the host.</p>	<p>NFS client service is required to be able to mount an NFS shared storage from where the Hyperscale Compliance Orchestrator will be able to read the source files and write the target files. For more information, see <a href="#">NFS Server Installation</a> (see page 138).</p>



Type	Host Requirement	Explanation
Hardware Requirements	<ul style="list-style-type: none"> <li>• Minimum: 8 vCPU, 64 GB of memory, 100GB data disk.</li> <li>• Recommended: 16 vCPU, 128GB of memory, 500GB data disk.</li> </ul>	OS disk space: 50 GB

### 5.5.3.2 Installation and setup (Podman Compose)



Delphix highly recommends new installations be performed on Kubernetes.

This section describes the steps you must perform to install the Hyperscale Compliance Orchestrator.

#### 5.5.3.2.1 Hyperscale Compliance installation

##### 5.5.3.2.1.1 Pre-requisites

Ensure that you meet the following requirements before you install the Hyperscale Compliance Orchestrator.

- Download the Hyperscale tar file ( `delphix-hyperscale-masking-15.0.0.tar.gz` ) from [download.delphix.com](https://download.delphix.com)<sup>46</sup>.
- You must create a user that has permission to install Podman and Podman Compose.
- Install Podman on VM. The minimum supported podman version is 4.4.1.

**Note:** By default, the Podman 4.4.1 version is not available for a few debian-based Linux distributions. For example, you cannot install Podman 4.4.1 on Ubuntu 20.04 and above. The workaround is to use a RHEL machine to host the hyperscale deployment with Podman.

- Install [Podman Compose](https://github.com/containers/podman-compose)<sup>47</sup> on the VM. The minimum supported podman-compose version is 1.0.6
- Check if podman and podman-compose are installed by running the following command:
  - `podman-compose -v` The above command displays an output similar to the following: `podman-compose version 1.0.6`
  - `podman -v` The above command displays an output similar to the following: `podman version 4.4.1`

<sup>46</sup> <https://download.delphix.com/folder/2224/Delphix%20Product%20Releases/Hyperscale%20Compliance>

<sup>47</sup> <https://github.com/containers/podman-compose>

- Podman can not create containers that bind to ports < 1024 ([here](#)<sup>48</sup>). Hyperscale's proxy container binds port 80, 443. Run following command to enable binding or port: `sudo sysctl net.ipv4.ip_unprivileged_port_start=80`
- [Only Required for Oracle Load Service] Download and install Linux-based [Oracle's instant client](#)<sup>49</sup> on the machine where the Hyperscale Compliance Orchestrator will be installed. The client should essentially include `instantclient-basic` (Oracle shared libraries) along with `instantclient-tools` containing Oracle's SQL\*Loader client. Both the packages `instantclient-basic` and `instantclient-tools` should be unzipped in the same directory. A group ownership id of 50 with a permission mode of 550 or a user id of 65436 with a permission mode of 500 must be set recursively on the directory where Oracle's instant client `binaries/libraries` will be installed. This is required by the Hyperscale Compliance Orchestrator to be able to read or execute from the directory.

 Oracle Load doesn't support Object Identifiers(OIDs).

#### 5.5.3.2.1.2 Procedure

Perform the following procedure to install the Hyperscale Compliance Orchestrator.

1. Unpack the Hyperscale tar file.

```
tar -xzf delphix-hyperscale-masking-15.0.0.tar.gz
```

2. Upon unpacking, you will find the podman image tar files which are categorized as below:

##### Universal images common for all connectors.

- `controller-service.tar`
- `masking-service.tar`
- `proxy.tar`

##### Oracle (required only for Oracle data source masking)

- `unload-service.tar`
- `load-service.tar`

##### MSSQL (required only for MS SQL data source masking)

- `mssql-unload-service.tar`
- `mssql-load-service.tar`

##### Delimited Files (required only for Delimited Files masking)

<sup>48</sup> <https://github.com/containers/podman/blob/main/rootless.md>

<sup>49</sup> <https://www.oracle.com/database/technologies/instant-client.html>

- `delimited-unload-service.tar`
- `delimited-load-service.tar`

#### **MongoDB (required only for MongoDB database masking)**

- `mongo-unload-service.tar`
- `mongo-load-service.tar`

#### **Parquet files (required only for Parquet file masking)**

- `parquet-unload-service.tar`
- `parquet-load-service.tar`

Each deployment set consists of 5 images (3 Universal images and 2 images related to each dataset type). Proceed to load the required images into podman as below:

For Oracle data source masking:

```
podman load --input unload-service.tar
podman load --input load-service.tar
podman load --input controller-service.tar
podman load --input masking-service.tar
podman load --input proxy.tar
```

For MS SQL data source masking:

```
podman load --input mssql-unload-service.tar
podman load --input mssql-load-service.tar
podman load --input controller-service.tar
podman load --input masking-service.tar
podman load --input proxy.tar
```

For Delimited Files masking:

```
podman load --input delimited-unload-service.tar
podman load --input delimited-load-service.tar
podman load --input controller-service.tar
podman load --input masking-service.tar
podman load --input proxy.tar
```


For MongoDB data source masking:

```
podman load --input mongo-unload-service.tar
podman load --input mongo-load-service.tar
podman load --input controller-service.tar
podman load --input masking-service.tar
podman load --input proxy.tar
```

For Parquet files masking:

```
podman load --input parquet-unload-service.tar
podman load --input parquet-load-service.tar
podman load --input controller-service.tar
podman load --input masking-service.tar
podman load --input proxy.tar
```

3. Create an NFS shared mount, that will act as a **Staging Area**, on the Hyperscale Compliance Orchestrator host where the Hyperscale Compliance Orchestrator will perform read/write/execute operations:
  - a. Create a 'Staging Area' directory. For example: `/mnt/hyperscale/staging_area`. The user(s) within each of the podman containers part of the Hyperscale Compliance Orchestrator and the appliance OS user(s) in the Continuous Compliance Engine(s), all have the user id as 65436 and/or group ownership id as 50. As such, the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, require the following permissions, based on the UID/GID of the OS user, so that the Hyperscale Compliance Orchestrator and the Continuous Compliance Engine(s) can perform read/write/execute operations on the staging area:
    - i. If the Hyperscale Compliance OS user has a UID of 65436, then the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, must have a UID of 65436 and 700 permission mode.
    - ii. If the Hyperscale Compliance OS user has a GID of 50 and does not have a UID of 65436, then the 'staging\_area' directory, along with the directory( `hyperscale` ) one level above, must have a GID of 50 and 770 permission mode.
  - b. Mount the NFS shared directory on the staging area directory( `/mnt/hyperscale/staging_area` ). This NFS shared storage can be created and mounted in two ways as detailed in the [NFS Server Installation](#)<sup>50</sup> section. Based on the umask value for the user which is used to mount, the permissions for the staging area directory could get altered after the NFS share has been mounted. In such cases, the permissions(i.e. 770 or 700 whichever applies based on point 3a) must be applied again on the staging area directory.

 The directory created in step 3a ('staging\_area') will be provided as the `mountName` and the corresponding shared path from the NFS file server as the `mountPath` in the MountFileSystems API.

4. Configure the following container volume bindings for the containers by editing the `podman-compose.yaml` file from tar:

<sup>50</sup> <https://portal.document360.io/continuous-compliance/docs/nfs-server-installation>

- a. For each of the containers, except the 'proxy' container, add a volume entry binding the staging area path (from 3(a), `/mnt/hyperscale`) to the Hyperscale Compliance Orchestrator container path(`/etc/hyperscale`) as a volume binding under the 'volumes' section.
- b. [Only Required for Oracle Load Service] For the **load-service** container, add a volume entry that binds the path of the directory on the host where both the Oracle instant Client packages were unzipped to the path on the container (`/usr/lib/instantclient`) under the 'volumes' section.
- c. [Only Required for Delimited Unload Service] For Delimited Files unload-service, the source NFS location has to be mounted to the container as volume in order for it to access the source files. The path mounted on the container is passed during the creation of the source connector-info.

```
# Mount your source NFS location onto your Hyperscale Engine server
sudo mount [-t nfs4] <source_nfs_endpoint>:<source_nfs_location>
<nfs_mount_on_host>
```

```
# Later mount <nfs_mount_on_host> as a podman volume to the delimited
unload-service container (in podman-compose.yaml, created using podman-
compose-delimitedfiles-sample.yaml)
unload-service:
  image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
  ...
  volumes:
    ...
    # Source files should be made available within the unload-service
    container file system
    # The paths within the container should be configured in the source
    section of connector-info [with type=FS]
    -
    <nfs_mount_on_host>:<source_files_mount_passed_during_connector_info_cre
    ation_path_in_contianer>
```

- d. [Only Required for Delimited load Service] For Delimited Files load-service, the target NFS location has to be mounted to the container as volume in order for it to access the target location where masked files will be placed. The path mounted on the container is passed during the creation of the target connector-info.

```
# Mount your target NFS location onto your Hyperscale Engine server
sudo mount [-t nfs4] <target_nfs_endpoint>:<target_nfs_location>
<target_nfs_mount_on_host>
```

```
# Later mount <nfs_mount_on_host> as a podman volume to the delimited
load-service container (in podman-compose.yaml, created using podman-
compose-delimitedfiles-sample.yaml)
load-service:
  image: delphix-delimited-load-service-app:${VERSION}
  ...
  volumes:
    ...
    # Target location should be made available within the load-service
container file system
    # The paths within the container should be configured in the target
section of connector-info [with type=FS]
    -
<target_nfs_mount_on_host>:<target_location_passed_during_connector_info
_creation_in_container>
```

5. [Optional] Some data (for example, logs, configuration files, etc.) that is generated inside containers may be useful to debug possible errors or exceptions while running the hyperscale jobs, and as such it may be beneficial to persist these logs outside containers. The following data can be persisted outside the containers:
  - a. The logs generated for each service i.e. unload, controller, masking, and load services.
  - b. The sqldr utility logs and control files at `opt/sqldr` location in the load-service container.
  - c. The file-upload folder at `/opt/delphix/uploads` in the controller-service container

If you would like to persist the above data on your host, then you have the option to do the same by setting up volume bindings in the respective service as indicated below, that map locations inside the containers to locations on the host in the `podman-compose.yaml` file. The host locations again must have a group ownership id of 50 with a permission mode of 770 or a user id of 65436 with a permission of 700, due to the same reasons as highlighted in step 3a.

Here are examples of the `podman-compose.yaml` file for Oracle, MS SQL, MongoDB data sources, Delimited file, and Parquet file masking:

For Oracle data source masking:

```
version: "4"
services:
  controller-service:
    image: delphix-controller-service-app:${VERSION}
    security_opt:
      - label:disable
    userns_mode: keep-id
    healthcheck:
      test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
      interval: 30s
      timeout: 25s
      retries: 3
      start_period: 30s
    depends_on:
```

```

- unload-service
- masking-service
- load-service
init: true
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-controller-data:/data
- /home/hyperscale_user/logs/controller_service:/opt/delphix/logs
- /mnt/hyperscale:/etc/hyperscale
environment:
- API_KEY_CREATE=${API_KEY_CREATE:-false}
- EXECUTION_STATUS_POLL_DURATION=${EXECUTION_STATUS_POLL_DURATION:-12000}
- LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_CONTROLLER_SERVICE:-INFO}
- API_VERSION_COMPATIBILITY_STRICT_CHECK=${
{API_VERSION_COMPATIBILITY_STRICT_CHECK:-false}
- LOAD_SERVICE_REQUIREPOSTLOAD=${LOAD_SERVICE_REQUIRE_POST_LOAD:-true}
- SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_UNLOAD_SPLIT_COUNT_VALIDATION:-false}
- SKIP_LOAD_SPLIT_COUNT_VALIDATION=${SKIP_LOAD_SPLIT_COUNT_VALIDATION:-false}
- CANCEL_STATUS_POLL_DURATION=${CANCEL_STATUS_POLL_DURATION:-60000}
unload-service:
image: delphix-unload-service-app:${VERSION}
security_opt:
- label:disable
usersns_mode: keep-id
init: true
environment:
- LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_UNLOAD_SERVICE:-INFO}
- UNLOAD_FETCH_ROWS=${UNLOAD_FETCH_ROWS:-10000}
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-unload-data:/data
- /mnt/hyperscale:/etc/hyperscale
- /home/hyperscale_user/logs/unload_service:/opt/delphix/logs
masking-service:
image: delphix-masking-service-app:${VERSION}
security_opt:
- label:disable
usersns_mode: keep-id
init: true
networks:
- hyperscale-net
restart: unless-stopped
volumes:
- hyperscale-masking-data:/data
- /mnt/hyperscale:/etc/hyperscale
- /home/hyperscale_user/logs/masking_service:/opt/delphix/logs
environment:
- LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_MASKING_SERVICE:-INFO}

```

```

- INTELLIGENT_LOADBALANCE_ENABLED=${INTELLIGENT_LOADBALANCE_ENABLED:-true}
load-service:
  image: delphix-load-service-app:${VERSION}
  security_opt:
    - label:disable
  userns_mode: keep-id
  init: true
  environment:
    - LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=${LOG_LEVEL_LOAD_SERVICE:-INFO}
    - SQLLDR_BLOB_CLOB_CHAR_LENGTH=${SQLLDR_BLOB_CLOB_CHAR_LENGTH:-20000}
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-load-data:/data
    - /mnt/hyperscale:/etc/hyperscale
    - /opt/oracle/instantclient_21_5:/usr/lib/instantclient
    - /home/hyperscale_user/logs/load_service:/opt/delphix/logs
    - /home/hyperscale_user/logs/load_service/sqlldr:/opt/sqlldr/
proxy:
  image: delphix-hyperscale-masking-proxy:${VERSION}
  init: true
  networks:
    - hyperscale-net
  ports:
    - "443:443"
  restart: unless-stopped
  depends_on:
    - controller-service
  #volumes:
  # Uncomment to bind mount /etc/config
  #- /nginx/config/path/on/host:/etc/config
networks:
  hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data:

```

For MS SQL data source masking: A sample file specific to MS SQL connector is should look like following.

```

version: "4"
services:
  controller-service:
    image: delphix-controller-service-app:${VERSION}
    security_opt:
      - label:disable
    userns_mode: keep-id
    healthcheck:

```



```

    test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
  interval: 30s
  timeout: 25s
  retries: 3
  start_period: 30s
  depends_on:
    - unload-service
    - masking-service
    - load-service
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-controller-data:/data
    - /home/hyperscale_user/logs/controller_service:/opt/delphix/logs
    - /mnt/hyperscale:/etc/hyperscale
  environment:
    - API_KEY_CREATE=${API_KEY_CREATE:-false}
    - EXECUTION_STATUS_POLL_DURATION=${EXECUTION_STATUS_POLL_DURATION:-12000}
    - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_CONTROLLER_SERVICE:-INFO}
    - API_VERSION_COMPATIBILITY_STRICT_CHECK=${
{API_VERSION_COMPATIBILITY_STRICT_CHECK:-false}
    - LOAD_SERVICE_REQUIREPOSTLOAD=${LOAD_SERVICE_REQUIRE_POST_LOAD:-true}
    - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=${
{SKIP_UNLOAD_SPLIT_COUNT_VALIDATION:-false}
    - SKIP_LOAD_SPLIT_COUNT_VALIDATION=${SKIP_LOAD_SPLIT_COUNT_VALIDATION:-false}
    - CANCEL_STATUS_POLL_DURATION=${CANCEL_STATUS_POLL_DURATION:-60000}
  unload-service:
    image: delphix-mssql-unload-service-app:${VERSION}
    security_opt:
      - label:disable
    userns_mode: keep-id
    init: true
    environment:
      - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_UNLOAD_SERVICE:-INFO}
      - UNLOAD_FETCH_ROWS=${UNLOAD_FETCH_ROWS:-10000}
      - SPARK_DATE_TIMESTAMP_FORMAT=${DATE_TIMESTAMP_FORMAT:-yyyy-MM-dd
HH:mm:ss.SSSS}
    networks:
      - hyperscale-net
    restart: unless-stopped
    volumes:
      - hyperscale-unload-data:/data
      - /mnt/hyperscale:/etc/hyperscale
      - /home/hyperscale_user/logs/unload_service:/opt/delphix/logs
  masking-service:
    image: delphix-masking-service-app:${VERSION}
    security_opt:
      - label:disable
    userns_mode: keep-id
    init: true

```

```

networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-masking-data:/data
  - /mnt/hyperscale:/etc/hyperscale
  - /home/hyperscale_user/logs/masking_service:/opt/delphix/logs
environment:
  - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_MASKING_SERVICE:-INFO}
  - INTELLIGENT_LOADBALANCE_ENABLED=${INTELLIGENT_LOADBALANCE_ENABLED:-true}
load-service:
  image: delphix-mssql-load-service-app:${VERSION}
  security_opt:
    - label:disable
  userns_mode: keep-id
  init: true
  environment:
    - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=${LOG_LEVEL_LOAD_SERVICE:-INFO}
    - SQLLDR_BLOB_CLOB_CHAR_LENGTH=${SQLLDR_BLOB_CLOB_CHAR_LENGTH:-20000}
    - SPARK_DATE_TIMESTAMP_FORMAT=${DATE_TIMESTAMP_FORMAT:-yyyy-MM-dd
HH:mm:ss.SSSS}
networks:
  - hyperscale-net
restart: unless-stopped
volumes:
  - hyperscale-load-data:/data
  - /mnt/hyperscale:/etc/hyperscale
  - /home/hyperscale_user/logs/load_service:/opt/delphix/logs
proxy:
  image: delphix-hyperscale-masking-proxy:${VERSION}
  init: true
networks:
  - hyperscale-net
ports:
  - "443:443"
restart: unless-stopped
depends_on:
  - controller-service
#volumes:
# Uncomment to bind mount /etc/config
#- /nginx/config/path/on/host:/etc/config
networks:
  hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data:

```

For Delimited Files masking: A sample file specific to the Delimited connector should look like the following.

```

version: "4"
services:
  controller-service:
    image: delphix-controller-service-app:<HYPERSCALE VERSION>
    security_opt:
      - label:disable
    userns_mode: keep-id
    healthcheck:
      test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
      interval: 30s
      timeout: 25s
      retries: 3
      start_period: 30s
    depends_on:
      - unload-service
      - masking-service
      - load-service
    init: true
    networks:
      - hyperscale-net
    restart: unless-stopped
    volumes:
      - hyperscale-controller-data:/data
      - /mnt/parent_staging_area:/etc/hyperscale
    environment:
      - API_KEY_CREATE=true
      - EXECUTION_STATUS_POLL_DURATION=120000
      - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
      - API_VERSION_COMPATIBILITY_STRICT_CHECK=false
      - LOAD_SERVICE_REQUIREPOSTLOAD=false
      - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=false
      - SKIP_LOAD_SPLIT_COUNT_VALIDATION=false
      - CANCEL_STATUS_POLL_DURATION=60000
      - SOURCE_KEY_FIELD_NAMES=unique_source_files_identifier
  unload-service:
    image: delphix-delimited-unload-service-app:<HYPERSCALE VERSION>
    security_opt:
      - label:disable
    userns_mode: keep-id
    init: true
    networks:
      - hyperscale-net
    restart: unless-stopped
    volumes:
      - hyperscale-unload-data:/data
      # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
      - /mnt/parent_staging_area:/etc/hyperscale
      # Source files should be made available within the unload-service container
file system

```

```

    # The paths within the container should be configured in the source section of
connector-info [with type=FS]
    - /mnt/source_files:/mnt/source
    #- /mnt/source_files2:/mnt/source2
masking-service:
    image: delphix-masking-service-app:<HYPERSCALE VERSION>
    security_opt:
        - label:disable
    userns_mode: keep-id
    init: true
    networks:
        - hyperscale-net
    restart: unless-stopped
    volumes:
        - hyperscale-masking-data:/data
        # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
        - /mnt/parent_staging_area:/etc/hyperscale
    environment:
        - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
        - INTELLIGENT_LOADBALANCE_ENABLED=true
load-service:
    image: delphix-delimited-load-service-app:<HYPERSCALE VERSION>
    security_opt:
        - label:disable
    userns_mode: keep-id
    init: true
    networks:
        - hyperscale-net
    restart: unless-stopped
    volumes:
        - hyperscale-load-data:/data
        # Staging area volume mount, here /mnt/parent_staging_area is used as an
example
        - /mnt/parent_staging_area:/etc/hyperscale
        # Target location should be made available within the load-service container
file system
        # The paths within the container should be configured in the target section of
connector-info [with type=FS]
        - /mnt/target_files:/mnt/target
        #- /mnt/target_files2:/mnt/target2
proxy:
    image: delphix-hyperscale-masking-proxy:<HYPERSCALE VERSION>
    init: true
    networks:
        - hyperscale-net
    ports:
        - "443:443"
        - "80:80"
    restart: unless-stopped
    depends_on:
        - controller-service
networks:

```

```

hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data:

```

For Parquet files masking:

```

version: "4"
services:
  controller-service:
    image: delphix-controller-service-app:${VERSION}
    healthcheck:
      test: 'curl --fail --silent http://localhost:8080/actuator/health | grep UP ||
exit 1'
      interval: 30s
      timeout: 25s
      retries: 3
      start_period: 30s
    depends_on:
      - unload-service
      - masking-service
      - load-service
    init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-controller-data:/data
    - /mnt/parent_staging_area:/etc/hyperscale
  environment:
    - API_KEY_CREATE=true
    - EXECUTION_STATUS_POLL_DURATION=120000
    - LOGGING_LEVEL_COM_DELPHIX_HYPERSCALE=INFO
    - API_VERSION_COMPATIBILITY_STRICT_CHECK=false
    - LOAD_SERVICE_REQUIREPOSTLOAD=false
    - SKIP_UNLOAD_SPLIT_COUNT_VALIDATION=false
    - SKIP_LOAD_SPLIT_COUNT_VALIDATION=false
    - CANCEL_STATUS_POLL_DURATION=60000
    - SOURCE_KEY_FIELD_NAMES=unique_source_files_identifier
  unload-service:
    image: delphix-parquet-unload-service-app:${VERSION}
    init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-unload-data:/data
    # Staging area volume mount, here /mnt/parent_staging_area is used as an example
    - /mnt/parent_staging_area:/etc/hyperscale
  environment:

```

```

- MAX_WORKER_THREADS_PER_JOB=512
# The default AWS region and credentials can be set using environment variables
#- AWS_DEFAULT_REGION=us-east-1
#- AWS_ACCESS_KEY_ID=<aws_access_key_id>
#- AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>
masking-service:
  image: delphix-masking-service-app:${VERSION}
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-masking-data:/data
    # Staging area volume mount, here /mnt/parent_staging_area is used as an example
    - /mnt/parent_staging_area:/etc/hyperscale
  environment:
    - LOGGING_LEVEL_COM_DELPPIX_HYPERSCALE=INFO
    - INTELLIGENT_LOADBALANCE_ENABLED=true
load-service:
  image: delphix-parquet-load-service-app:${VERSION}
  init: true
  networks:
    - hyperscale-net
  restart: unless-stopped
  volumes:
    - hyperscale-load-data:/data
    # Staging area volume mount, here /mnt/parent_staging_area is used as an example
    - /mnt/parent_staging_area:/etc/hyperscale
  #environment:
    # The default AWS region and credentials can be set using environment variables
    #- AWS_DEFAULT_REGION=us-east-1
    #- AWS_ACCESS_KEY_ID=<aws_access_key_id>
    #- AWS_SECRET_ACCESS_KEY=<aws_secret_access_key>
proxy:
  image: delphix-hyperscale-masking-proxy:${VERSION}
  init: true
  networks:
    - hyperscale-net
  ports:
    - "443:443"
    - "80:80"
  restart: unless-stopped
  depends_on:
    - controller-service
networks:
  hyperscale-net:
volumes:
  hyperscale-load-data:
  hyperscale-unload-data:
  hyperscale-masking-data:
  hyperscale-controller-data

```

6. (OPTIONAL) To modify the default Hyperscale configuration properties for the application, see [Configuration Settings](#)<sup>51</sup>.

7. Run the application from the same location where you extracted the `podman-compose.yaml` file.


```
podman-compose up -d
```

- Run the following command to check if the application is running. The output of this command should show five containers up and running.

```
podman-compose ps
```

- Run the following command to access application logs of a given container.

```
podman logs -f service_container_name>
```


 Service container name can be accessed by output of the command `podman-compose ps`.

- Run the following command to stop the application (if required).

```
podman-compose down
```

8. Once the application starts, an API key will be generated that will be required to authenticate with the Hyperscale Compliance Orchestrator. This key will be found in the podman container logs of the controller service. You can either look for the key from the controller service logs location that was set as a volume binding in the `podman-compose.yaml` file or you could use the following 'podman' command to retrieve the logs.

```
podman logs -f <service_container_name>
```

 Service container name can be accessed by output of the command `podman-compose ps`.

The above command displays an output similar to the following where the string `NEWLY GENERATED API KEY` can be grepped from the log::

```
2022-05-18 12:24:10.981 INFO 7 --- [          main] o.a.c.c.C.[Tomcat].[localhost].
[/] : Initializing Spring embedded WebApplicationContext
2022-05-18 12:24:10.982 INFO 7 --- [          main]
w.s.c.ServletWebServerApplicationContext : Root WebApplicationContext: initialization
completed in 9699 ms
```

<sup>51</sup> <https://portal.document360.io/continuous-compliance/docs/configuration-settings-5>

```
NEWLY GENERATED API KEY: 1.89\PH1dH5JQwHuQvzawD99sf4SpBPXJADUmJS8v00VCF4V7rjtRFAftGWy
gFfsqM
```

To authenticate with the Hyperscale Compliance Orchestrator, you must use the API key and include the HTTP Authorization request header with the type apk; apk <API Key> .

For more information, see the **Authentication** section under [Accessing the Hyperscale Compliance API](#) (see [page 141](#)).

### 5.5.3.2 Continuous Compliance Engine Installation

Delphix Continuous Compliance Engine is a multi-user, browser-based web application that provides complete, secure, and scalable software for your sensitive data discovery, masking, and tokenization needs while meeting enterprise-class infrastructure requirements. For information about installing the Continuous Compliance Engine, see [Continuous Compliance Engine Installation](#)<sup>52</sup> documentation.

### 5.5.3.3 How to generate a support bundle (Podman Compose)



If it is not already installed, you must install bash shell and yq to generate a Hyperscale support bundle. For more information on downloading yq, refer to the [yq Downloads page on GitHub](#)<sup>53</sup>.

#### 5.5.3.3.1 1. Find the “generate\_support\_bundle.sh” script

- Login to Hyperscale VM for which you want to generate the support bundle.
- generate\_support\_bundle.sh” file is bundled with the release tar file. You can find this script under `tools/support-scripts` folder, (present under the directory, where you will untar the release tar file on Hyperscale Engine). For example, `/path_to_untarred_hyperscale_product/tools/support-scripts` .

#### Example:

```
dlpxuser@delphix:~/test$ cd tools/support-scripts/
dlpxuser@delphix:~/test$ ls -ltr
total 48
-rwxr-xr-x  1 delphix  staff   823 Jul  7 09:55 generate_support_bundle.sh
-rwxr-xr-x  1 delphix  staff   463 Jul  7 09:55 container_information.sh
-rwxr-xr-x  1 delphix  staff  5597 Jul  7 09:55 collect_container_support_info.sh
-rw-r--r--  1 delphix  staff  5316 Jul  7 09:55 README.md
```

<sup>52</sup> <https://portal.document360.io/continuous-compliance/docs/first-time-setup>

<sup>53</sup> [https://github.com/mikefarah/yq/#wf\\_ssl\\_install](https://github.com/mikefarah/yq/#wf_ssl_install)



### 5.5.3.3.2 2. Modify the “container\_information.sh” script parameters

- Uncomment following 3 lines in file.

```
shopt -s expand_aliases
alias docker=podman
alias docker-compose=podman-compose
```

- Change the mount\_path and docker\_compose\_file\_path accordingly.

**Example:**

```
mount_path=/home/delphix/hyperscale
docker_compose_file_path=/home/delphix/podman-compose.yaml
```



- mount\_path: Absolute path configured for mount directory in podman-compose file which is mapped to /etc/hyperscale .
- podman\_compose\_file\_path: Absolute path for podman\_compose.yaml file.

### 5.5.3.3.3 3. Execute the “generate\_support\_bundle.sh” script

- Execute the “generate\_support\_bundle.sh” script from tools/support-scripts/ folder.

**Example:**

```
d1pxuser@delphix:~/test/tools/support-scripts/$ ./generate_support_bundle.sh
....
Generating support bundle tar file...
....
```

- Enter the “Password” when prompted.

### 5.5.3.3.4 4. Find the Generated Support Bundle Tar File

The resulting support bundle will be located at /etc/hyperscale/hyperscale-support-\*\*\*\*.tar.gz inside the container. This means the tar file is generated under the path which is mapped to /etc/hyperscale in podman-compose file and is directly accessible from Hyperscale VM.

**Example:**

```
dlpxuser@delphix:~/test$ ls -ltr ../hyperscale/
total 316
drwxrwxrwx 5 1004 1005 4096 Feb  9 10:14 aks-mount
-rw-r--r-- 1 65436 staff 104189 Feb 17 08:52 hyperscale-support-
<current_timestamp>.tar.gz
```

**The support bundle tar file contains the following information:**

- Hyperscale Logs
- The output of mpstat for CPU utilization info.
- The output of proc/meminfo for memory info.
- The output of proc/cpuinfo for cpu info.
- Files to show the memory limit for the application container and the max usage of the app container in bytes.
- Redacted database file to restore the Hyperscale VM
- podman compose file



- The script generate\_support\_bundle.sh generates a bare-bones support bundle from a Hyperscale engine running in podman.
- Execute the generate\_support\_bundle.sh from the untar location.
- The resulting support bundle will be at /etc/hyperscale/hyperscale-support-\*\*\*\*.tar.gz inside the container. This means the tar file is generated under a path that is mapped to /etc/hyperscale in podman-compose file and is directly accessible from Hyperscale VM.
- The user should have privileges or permission to execute the podman command in order to generate the support bundle.

## 5.6 NFS server installation

The Hyperscale Compliance Orchestrator requires a Staging Area to read from the source file(s) and write to the target file(s). The Staging Area must be an NFS-shared filesystem accessible to the Hyperscale Compliance Orchestrator and the Continuous Compliance Engines. The following are the supported ways by which the filesystem can be shared over NFS(NFSv3/NFSv4):

### 5.6.1 Delphix Continuous Data Engine empty VDB

To create a Delphix Virtualization Engine empty VDB, follow the below procedure.

### 5.6.1.1 Continuous Data Engine installation

Delphix Virtualization Engine is a data management platform that provides the ability to securely copy and share datasets. Using virtualization, you will ingest your data sources and create virtual data copies, which are full read-write capable database instances that use a small fraction of the resources a normal database copy would require.

For information about installing the Virtualization Engine, see [Virtualization Engine Installation](#)<sup>54</sup> documentation.

### 5.6.1.2 Discover and configure Hyperscale Compliance Orchestrator environment

1. After installing and configuring the Virtualization Engine, make sure that the [Network and Connectivity Requirements](#)<sup>55</sup> for using Empty VDB on Unix environments are met.
2. Discover the Hyperscale Compliance Orchestrator Unix host on the Virtualization's Engine Management application. For more information, see [Adding a Unix Environment](#)<sup>56</sup>.
3. Navigate to **Manage > Environments** to view the discovered Hyperscale Compliance Orchestrator Unix host.
4. After the discovery is completed, configure the same Unix host on the Environments screen such that the IP addresses of the Hyperscale Compliance Orchestrator Unix host along with the Continuous Compliance Engines part of the Continuous Compliance Engine cluster are populated in the NFS Addresses field. This is done to ensure that the empty VDB is shared with both Hyperscale Compliance Orchestrator and the Continuous Compliance Engines part of the Continuous Compliance Engine cluster.

---

<sup>54</sup> <https://docs.delphix.com/docs/deployment>

<sup>55</sup> <https://docs.delphix.com/docs/datasets/unstructured-files-and-app-data/unstructured-files-environment-requirements/unstructured-files-on-unix-environments/network-and-connectivity-requirements-for-unix-environments>

<sup>56</sup> <https://docs.delphix.com/docs/datasets/unstructured-files-and-app-data/unstructured-files-environment-requirements/unstructured-files-on-unix-environments/adding-a-unix-environment>

The screenshot displays the Delphix Management console for a 'Hyperscale Compliance Engine Unix Host'. The interface is divided into several sections:

- Navigation:** 'Manage Resources System Help' at the top.
- Left Panel:** 'Environments' with a filter set to 'none' and a list of environments including 'Hyperscale Compliance En... Unix Host - Enabled'.
- Main Content:**
  - Details Tab:** Shows 'Basic Information' and 'ENVIRONMENT USERS'. The users table has one entry: 'delphix' with a star icon.
  - Attributes:**
    - Host Address: 10.119.138.38
    - NFS Addresses: 10.119.138.38, 10.119.138.23, 10.119.138.24, 10.119.138.25, 10.119.138.26, 10.119.138.29, 10.119.138.31
    - SSH Port: 22
    - DSP KeyStore Path: None
    - DSP KeyStore Alias: None
  - Listeners:** Section for configuring listeners.
  - Java Development Kit:** Shows 'Installed Java Development Kit' and 'Default JDK'.
  - FAULTS:** A large empty box with the text 'No Faults'.

### 5.6.1.3 Provision an empty VDB

1. Follow the steps listed under [Create an Empty VDB for Unstructured Files in the Delphix Engine](#)<sup>57</sup> to provision an empty VDB on the discovered Hyperscale Compliance Orchestrator Unix host.
2. Note the mount path provided while provisioning the empty VDB as that is the path which will be used to fill the empty VDB with the source file(s) that the Hyperscale Compliance Orchestrator needs to mask and where the target masked file(s) will be placed.

**i** Hyperscale Compliance OS user should have read/write permissions on the mount point path where the empty VDB will be provisioned. Hyperscale Compliance OS user should have read/write permissions on the mount point path where the empty VDB will be provisioned.

The location of the mounted empty VDB on the Hyperscale Compliance Orchestrator Unix host can be found with a simple 'grep' of the mount path, provided while provisioning the empty VDB, using the 'mount' utility:

```
hyperscale-engine:~$ df -h | grep /mnt/provision/hyperscale_data
10.119.138.34:/domain0/group-2/appdata_container-3/appdata_timeflow-4/datafile 20T 3.5T
```

<sup>57</sup> <https://docs.delphix.com/docs/datasets/unstructured-files-and-app-data/create-an-empty-vdb-for-unstructured-files-in-the-delphix-engine>

```
16T 18% /mnt/provision/hyperscale_data
```

3. Copy the source file(s) to the location where the empty VDB has been mounted.

## 5.6.2 NFS file server

1. An NFS shared filesystem can also be provided by a typical NFS server. Export a filesystem from the NFS file server such that the Hyperscale Compliance Orchestrator and Continuous Compliance Engines part of the Continuous Compliance Engine Cluster have read and write permission on it. As such, the export entry should be of the following form based on the UID/GID corresponding to the owner of the shared path:

```
<mount_path> <ip1,ip2,ip3,ipn>(rw,all_squash,anonuid=<uid>,anongid=<gid>)
```

2. Export the NFS share using the below command:

```
sudo exportfs -rav
```

3. Once the NFS share is exported from the NFS server, proceed to mount the same share on the Hyperscale Compliance Orchestrator host:

```
sudo mount -t nfs -o vers=4 <nfs-server-host-ip>:<mount_path>
<user.home>/hyperscale/mount-dir
```

## 5.6.3 Storage requirements for the NFS file server

Considering a single Hyperscale Compliance job execution, the Hyperscale Compliance Orchestrator will store unloaded files (unloaded from source) and masked files. As such, the required storage will amount to 2X the size of the source data.

## 5.7 Accessing the Hyperscale Compliance API

Open a web browser and type the following in the address bar: `https://<hyperscale-compliance-host-address>/hyperscale-compliance`. Replace the `hyperscale-compliance-host-address` with the IP address of the Hyperscale Compliance Orchestrator VM.

Sample Command to determine the IP address:

```
# Get your hyperscale masking proxy service Cluster IP
# Sample command for LINUX
kubectl describe service proxy -n hyperscale-services |grep "IP:" | tr -s " " | cut
-d " " -f 2
```

## 5.7.1 Authentication

To authenticate with the Hyperscale Compliance Orchestrator, you must use an API key. It is done by including the key in the HTTP Authorization request header with the type apk.

An example cURL command with the API Key looks like the following:

```
curl --header 'Authorization: apk
1.t8YTjLyPiMatdtnhAw9RD0gRVZr2hFsrfikp3YxVl8URdB9zuaVHcMuhXkLd1TLj'
```

As described in the [HTTP Authorization request header](#)<sup>58</sup> documentation, the following is the typical syntax for the authorization header:


```
Authorization: <auth-scheme> <authorisation-parameters>
```

For Basic Authentication, You must include the following header parameters: `Authorization: Basic <credentials>`

For the Bearer Authentication scheme, you must use the following: `Authorization: Bearer <JWT Bearer Token>`

## 5.7.2 Creating an API key

An API key is a simple encrypted string that you can use when calling Hyperscale Compliance APIs.

 You must use the initial created API key to create a new secure key. It is done by creating a new API Client entity. The “name” attribute must be the desired name to uniquely identify the user of this key. For more information about initial created API key, refer to step 8 under the [Generate a New Key](#)<sup>59</sup> section.

Run the following command to create a new API key.

```
curl -X 'POST' \
  'https://<host-name>/api/<api_version>/api-keys' \
  -H 'accept: application/json' \
  -H 'Authorization: apk
1.t8YTjLyPiMatdtnhAw9RD0gRVZr2hFsrfikp3YxVl8URdB9zuaVHcMuhXkLd1TLj' \
  -H 'Content-Type: application/json' \
  -d '{
  "name": "<name-of-key>"
}'
```

<sup>58</sup> <https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Authorization>

<sup>59</sup> <https://delphixdocs.atlassian.net/hyperscale-compliance-7-0-0/docs/installation-2>

The above command displays a response message similar to the following:

**i** Copy or save the newly created token from the response as this token value will not be accessible later.

```
{
  "api_key_id": 2,
  "token": "2.ExZtmf6EN1xvFMsXpXl0yhHVYlTuFzCm2yGhpU0QQ5ID8N8oGz79d4yn8ZsPhF46"
}
```

Since you have created a new and secure API key, you must delete the old key for security reasons.

Run the following command to delete the old key.

```
curl -X 'DELETE' \
  'https://<host-name>/api/<api_version>/api-keys/1' \
  -H 'accept: */*' \
  -H 'Authorization: apk
2.ExZtmf6EN1xvFMsXpXl0yhHVYlTuFzCm2yGhpU0QQ5ID8N8oGz79d4yn8ZsPhF46'
```

### 5.7.3 Using the newly generated key

After you delete the old key, revert the changes performed in step 5 of the [Hyperscale Compliance Installation](#) (see page 76) and restart docker-compose.

You must be able to use the new key for authorization as follows:

```
curl --header 'Authorization: apk
2.ExZtmf6EN1xvFMsXpXl0yhHVYlTuFzCm2yGhpU0QQ5ID8N8oGz79d4yn8ZsPhF46'
```

### 5.7.4 Default API Version

If the version is omitted from the base path of the request's URL, a default API version i.e. the latest API version of that Hyperscale Engine is used.

## 6 How to setup a Hyperscale Compliance job

### 6.1 Pre-checks

You must check the following before starting a job:

- Storage space must be 2 times the size of the source data for NFS storage.
- You must have sufficient storage in the target DB for loading the masked data.
- You must check and increase the size of the temporary tablespace in Oracle. For example, if you have 4 billion rows, then you must use 100G.
- You must check and provide the required permission(i.e. 770 or 700) after creating an empty VDB(or mounting an NFS share) on the mount folder on the Hyperscale Compliance host.
- Based on the unmask value for the user that is used to mount, the permissions for the staging area directory could get altered after the empty VDB or NFS share has been mounted. In such cases, you must re-apply the permissions (i.e. 770 or 700) on the staging area directory.
- You must restart the services after changing the permission on VDB mounted folder in case you already have created the containers.
- Continuous Compliance Engine should be cleaned up before use and should only be used with Hyperscale Job. Any other masking job on Continuous Compliance Engine apart from Hyperscale Compliance Orchestrator will impact the performance of Hyperscale Compliance jobs.
- Currently, the Hyperscale Compliance Orchestrator doesn't provide the ability to allow you to configure the masking job behavior in case of non-conformant data and does not process non-conformant data warnings from the Delphix Continuous Compliance Engine. Therefore, it is recommended to verify the value of `DefaultNonConformantDataHandling` algorithm group setting on all the Hyperscale Compliance Orchestrator. For more information, refer to the [Algorithm Group Settings](#)<sup>60</sup> section. It is recommended to set the value to FAIL so that Hyperscale Job will also fail instead of leaving the data unmasked.
- If the table that you are masking has a column type of BLOB/CLOB, then you must have a minimum of 2GB memory per CLOB/BLOB column. Depending upon the unload-split you are using, you may need to increase this memory in multiple of that. For example, if you have 4 tables (each with 1 column as BLOB/CLOB type) and unload-split is 3, then your memory requirement on the Hyperscale Compliance host will be:  $(4(\text{no. of tables}) \times 2(\text{memory required per CLOB/BLOB column}) \times 3(\text{unload-split used}))\text{GB} + 16 \text{ GB (minimum required memory for running Hyperscale Compliance Orchestrator)} = 40 \text{ GB approx.}$

### 6.2 API Flow to Setup a Hyperscale Compliance Job

The following is the API flow for setting up and executing a Hyperscale Compliance job.


---

<sup>60</sup> <https://portal.document360.io/continuous-compliance/docs/masking-api-client#algorithm-group-settings>



1. Register Continuous Compliance Engine(s)
2. Create a Mount Point
3. Create Connector Info
4. Upload format file using `POST /file-upload` endpoint. For more information, refer to [Hyperscale Compliance API \(see page 200\)](#). [Required only for those who need to mask embedded XML/JSON data]
5. Create Structured Data Format [Required only for those who need to mask embedded XML/JSON data]
6. Create a Dataset
7. Create a Job
8. Create Execution

The following are the sample API requests/responses for a typical Hyperscale Compliance job execution workflow. The APIs can be accessed using a swagger-based API client by accessing the following URL; `https://<hyperscale-compliance-host-address>/hyperscale-compliance`.

 APIs must be called only in the below order.

## 6.3 Engines API

**POST /engines (register an engine):**

**Request:**

```
{
  "name": "Delphix Continuous Compliance Engine 6.0.14.0 on AWS",
  "type": "MASKING",
  "protocol": "http",
  "hostname": "de-6014-continuous-compliance.delphix.com",
  "username": "hyperscale_compliance_user",
  "password": "password123"
}
```

**Response:**

```
{
  "id": 1,
  "name": "Delphix Continuous Compliance Engine 6.0.14.0 on AWS",
  "type": "MASKING",
  "protocol": "http",
  "hostname": "de-6014-continuous-compliance.delphix.com",
```

```

"username": "hyperscale_compliance_user",
"ssl": true,
"ssl_hostname_check": true
}

```

## 6.4 MountFileSystems API

### POST /mount-fileSystems (create a file mount)

#### Request:

```

{
"mountName": "staging_area",
"hostAddress": "de-6014-continuous-data.dlpxdc.co",
"mountPath": "/domain0/group-2/appdata_container-12/appdata_timeflow-13/datafile",
"mountType": "NFS4",
"options": "rw"
}

```

#### Response:

```

{
"id": 1,
"mountName": "staging_area",
"hostAddress": "de-6014-continuous-data.dlpxdc.co",
"mountPath": "/domain0/group-2/appdata_container-12/appdata_timeflow-13/datafile",
"mountType": "NFS4",
"options": "rw"
}

```

## 6.5 ConnectorInfo API

### POST /connector-info (create connector info for Hyperscale Compliance)

#### Oracle Request:

```

{
"source": {
"jdbc_url": "jdbc:oracle:thin:@oracle-19-src.dlpxdc.co:1521/VDBOMSRDC20SRC",
"user": "oracle_db_user",
"password": "password123"
},
"target": {
"jdbc_url": "jdbc:oracle:thin:@rh79-ora-19-tgt.dlpxdc.co:1521/VDBOMSRDC200B_TGT",
"user": "oracle_db_user",
"password": "password123"
}
}

```

```
}
}
```

**Oracle Response:**

```
{
  "id": 1,
  "source": {
    "jdbc_url": "jdbc:oracle:thin:@oracle-19-src.dlpxdc.co:1521/VDBOMSRDC20SRC",
    "user": "oracle_db_user"
  },
  "target": {
    "jdbc_url": "jdbc:oracle:thin:@rh79-ora-19-tgt.dlpxdc.co:1521/VDBOMSRDC200B_TGT",
    "user": "oracle_db_user"
  }
}
```

**Example 2: This example is for the cases where either username or password needs to be in either uppercase or camel case**

**Oracle Request:**

```
{
  "source": {
    "user": "\"y2ijf0oj2\"",
    "password": "\"xyz\"",
    "jdbc_url": "jdbc:oracle:thin:@xyz.com:1521/DBOMSRDC200B",
    "connection_properties": {}
  },
  "target": {
    "jdbc_url": "jdbc:oracle:thin:@xyx.com:1521/DBOMSRDC200B",
    "user": "\"y2ijf0oj2\"",
    "password": "\"xyz\"",
    "connection_properties": {}
  }
}
```

**Oracle Response:**

```
{
  "id": 1,
  "source": {
    "user": "\"y2ijf0oj2\"",
    "jdbc_url": "jdbc:oracle:thin:@xyz.com:1521/DBOMSRDC200B",
    "connection_properties": {}
  },
  "target": {
    "jdbc_url": "jdbc:oracle:thin:@xyx.com:1521/DBOMSRDC200B",
    "user": "\"y2ijf0oj2\"",
    "connection_properties": {}
  }
}
```

```
}
}
```

**MSSQL Request:**

```
{
  "source": {
    "jdbc_url": "jdbc:sqlserver://hyperscale-
mssql.dlpxdc.co;database=SourceDB2019;instanceName=SQL2019",
    "user": "sa",
    "password": "password123"
  },
  "target": {
    "jdbc_url": "jdbc:sqlserver://hyperscale-
mssql.dlpxdc.co;database=SourceDB2019;instanceName=SQL2019;",
    "user": "sa",
    "password": "password123"
  }
}
```

**MSSQL Response:**

```
{
  "id": 1,
  "source": {
    "user": "sa",
    "jdbc_url": "jdbc:sqlserver://hyperscale-
mssql.dlpxdc.co;database=SourceDB2019;instanceName=SQL2019"
  },
  "target": {
    "jdbc_url": "jdbc:sqlserver://hyperscale-
mssql.dlpxdc.co;database=SourceDB2019;instanceName=SQL2019;",
    "user": "sa",
  }
}
```

**Delimited Files Request:**

```
{
  "source": {
    "type": "FS",
    "properties": {
      "server": "local",
      "path": "/mnt/source"
    }
  },
  "target": {
    "type": "FS",
    "properties": {
      "server": "local",

```

```

    "path": "/mnt/target"
  }
}

```

**i** Here `type=FS` means **“File System”**. Currently, the Delimited Files connector only supports files mounted directly onto the docker container, i.e. available on the container file system. In the above example `/mnt/source` & `/mnt/target`` are paths inside the container that denote the source and target location respectively.

#### Delimited Files Response:

```

{
  "id": 1,
  "source": {
    "type": "FS",
    "properties": {
      "server": "local",
      "path": "/mnt/source"
    }
  },
  "target": {
    "type": "FS",
    "properties": {
      "server": "local",
      "path": "/mnt/target"
    }
  }
}

```

#### MongoDB Connector Request:

```

{
  "source": {
    "mongo_url": "mongodb://<hostname>:<port>",
    "user": "mongo_user",
    "password": "mongo_password"
  },
  "target": {
    "mongo_url": "mongodb://<hostname>:<port>/?
replicaSet=<mongo_rs>&tls=true&tlsCertificateKeyFile=<cert_path>",
    "user": "mongo_user",
    "password": "mongo_password"
  }
}

```

**MongoDB Connector Response:**

```
{
  "source": {
    "mongo_url": "mongodb://mongodb-src.example:27017",
    "user": "mongo_user",
  },
  "target": {
    "mongo_url": "mongodb://mongodb-tgt.example.com:27017",
    "user": "mongo_user",
  }
}
```



A failure in the load may leave the target datasource in an inconsistent state since the load step truncates the target when it begins. If the source and target data source are configured to be the same datasource and a failure occurs in the load step, it is recommended that the single datasource be restored from a backup (or use the continuous data engine's rewind feature if you have a VDB as the single datasource) after the failure in the load step as the datasource may be in an inconsistent state. After the datasource is restored, you may kick off another hyperscale job. If the source and target data source are configured to be different, you may use the Hyperscale Compliance Orchestrator restartability feature to restart the job from the point of failure in the load/post-load step.

**Parquet Request:****Example1 - Without AWS S3 bucket credentials:**

```
{
  "source": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_source/sub_folder"
    }
  },
  "target": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_target/sub_folder"
    }
  }
}
```

**Example2 - With AWS S3 bucket credentials:**

```
{
  "source": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_source/sub_folder",
      "aws_region": "us-east-1",
      "aws_access_key_id": "AKIAYHUJKLDHMB",
      "aws_secret_access_key": "x2IXoHDYHhdydmmm&h12563kaka",
      "aws_role_arn": "56436882398"
    }
  },
  "target": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_target/sub_folder",
      "aws_region": "us-east-1",
      "aws_access_key_id": "AKIAYHUJKLDHMB",
      "aws_secret_access_key": "x2IXoHDYHhdydmmm&h12563kaka",
      "aws_role_arn": "56436882398"
    }
  }
}
```

Overview of the parameters:

- **type:** Here we refer to the type of connector we are creating. Currently, we only support “AWS” which refers to “Cloud vendor”, indicating that all the source and target files are available within the container through AWS credentials.
- **properties:** Holds the server and path values.
- **server:** Points to the S3 location. Currently, this is ignored as the only supported type is “AWS”.
- **path:** The path to the AWS source/target S3 bucket location. This path needs not be the exact path to the files, but a parent directory. In case you are planning to use a profiler then this path must be the exact path to the parquet files.
- Additional parameters [in case AWS credentials need to be passed for separately]:
  - **aws\_region:** The AWS region the S3 bucket is part of.
  - **aws\_access\_key\_id:** The AWS access key ID generated for the AWS Role.
  - **aws\_secret\_access\_key:** The AWS secret access key generated for the AWS Role.
  - **aws\_role\_arn:** The AWS Role Identifier.

**Parquet Response:**

**For Example1**

```
{
  "source": {
    "type": "AWS",
    "properties": {
```

```

    "server": "S3",
    "path": "s3_bucket_source/sub_folder"
  }
},
"target": {
  "type": "AWS",
  "properties": {
    "server": "S3",
    "path": "s3_bucket_target/sub_folder"
  }
}
}
}

```

**For Example2**

```

{
  "source": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_source/sub_folder",
      "aws_region": "us-east-1",
      "aws_access_key_id": "AKIA*****",
      "aws_secret_access_key": "x2IX*****",
      "aws_role_arn": "56436882398"
    }
  },
  "target": {
    "type": "AWS",
    "properties": {
      "server": "S3",
      "path": "s3_bucket_target/sub_folder",
      "aws_region": "us-east-1",
      "aws_access_key_id": "AKIA*****",
      "aws_secret_access_key": "x2IX*****",
      "aws_role_arn": "56436882398"
    }
  }
}
}

```

## 6.6 StructuredDataFormat APIs

This functionality is specifically needed when there is a requirement to mask XML/JSON data contained within database columns. These endpoints require a `file_upload_ref` that can be generated via the `POST /file-upload` endpoint wherein you must upload the format of the XML/JSON data to be masked. For more information, refer to the [Hyperscale Compliance API \(see page 200\)](#) documentation.

`POST /structured-data-format` (create structured-data-format for a column)



**Request:**

```
{
  "file_upload_ref": "delphix-file://upload/f_XXXX/XXX.xml",
  "doc_type": "XML",
  "masking_inventory_paths": [
    {
      "path": "/catalog/book/author",
      "domain_name": "NULL_SL",
      "algorithm_name": "dlpx-core:FullName"
    },
    {
      "path": "/catalog/book/isbn",
      "domain_name": "ADDRESS",
      "algorithm_name": "dlpx-core:CM Alpha-Numeric"
    }
  ]
}
```

**Response:**

```
{
  "structured_data_format_id": 1,
  "file_upload_ref": "delphix-file://upload/f_XXXX/XXX.xml",
  "doc_type": "XML",
  "masking_inventory_paths": [
    {
      "path": "/catalog/book/author",
      "domain_name": "NULL_SL",
      "algorithm_name": "dlpx-core:FullName"
    },
    {
      "path": "/catalog/book/isbn",
      "domain_name": "ADDRESS",
      "algorithm_name": "dlpx-core:CM Alpha-Numeric"
    }
  ]
}
```

## 6.7 DataSets API



- Table and schema names are case-sensitive.

- For the MSSQL connector, it's recommended to provide filter\_key if the unload\_split count is more than 1, and the table does not contain a primary/unique key, else data will be unloaded through a sequential approach which may be slower. If filter\_key is not provided and the table includes a primary/unique key then Hyperscale will scan the table to fetch the key automatically.
- The dataset date format should be the same as the environment variable date format value.
- In one dataset, all the inventories should use the same date format for all date formats.

## POST /data-sets (create dataset for Hyperscale Compliance)

**i** Alternatively, you can create or update the dataset using payload in a file with the below endpoints:

- `POST /data-sets/file-upload`
- `PUT /data-sets/file-upload/{dataSetId}`

The above endpoints require a `file_upload_ref` that can be generated via the `POST /file-upload` endpoint. For more information, refer to the [Hyperscale Compliance API \(see page 200\)](#) documentation.

## Request (with single table):

```
{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "schema_name": "SCHEMA_1",
        "table_name": "TABLE_1",
        "unload_split": 4
      },
      "target": {
        "schema_name": "SCHEMA_1_TARGET",
        "table_name": "TABLE_1_TARGET",
        "stream_size": 65536
      },
      "masking_inventory": [
        {
          "field_name": "FIRST_NAME",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "FirstNameLookup"
        },
        {
          "field_name": "LAST_NAME",
```

```

    "domain_name": "LAST_NAME",
    "algorithm_name": "LastNameLookup"
  },
  {
    "field_name": "XmlData",
    "structured_data_format_id": 1
  }
]
}
]
}

```

**Response (with single table):**

```

{
  "id": 1,
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "schema_name": "SCHEMA_1",
        "table_name": "TABLE_1",
        "unload_split": 4
      },
      "target": {
        "schema_name": "SCHEMA_1",
        "table_name": "TABLE_1",
        "stream_size": 65536
      },
      "masking_inventory": [
        {
          "field_name": "FIRST_NAME",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "FirstNameLookup"
        },
        {
          "field_name": "LAST_NAME",
          "domain_name": "LAST_NAME",
          "algorithm_name": "LastNameLookup"
        },
        {
          "field_name": "XmlData",
          "structured_data_format_id": 1
        }
      ]
    }
  ]
}
]
}

```

**Request (with single table & filter\_key in the source):**

```

{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "schema_name": "SCHEMA_1",
        "table_name": "TABLE_1",
        "unload_split": 4,
        "filter_key": "PKID"
      },
      "target": {
        "schema_name": "SCHEMA_1_TARGET",
        "table_name": "TABLE_1_TARGET",
        "stream_size": 65536
      },
      "masking_inventory": [
        {
          "field_name": "FIRST_NAME",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "FirstNameLookup"
        },
        {
          "field_name": "LAST_NAME",
          "domain_name": "LAST_NAME",
          "algorithm_name": "LastNameLookup"
        },
        {
          "field_name": "XmlData",
          "structured_data_format_id": 1
        }
      ]
    }
  ]
}

```

**Response (with single table & filter\_key in the source):**

```

{
  "id": 1,
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "schema_name": "SCHEMA_1",
        "table_name": "TABLE_1",
        "unload_split": 4,
        "filter_key": "PKID"
      },
      "target": {

```

```

    "schema_name": "SCHEMA_1",
    "table_name": "TABLE_1",
    "stream_size": 65536
  },
  "masking_inventory": [
    {
      "field_name": "FIRST_NAME",
      "domain_name": "FIRST_NAME",
      "algorithm_name": "FirstNameLookup"
    },
    {
      "field_name": "LAST_NAME",
      "domain_name": "LAST_NAME",
      "algorithm_name": "LastNameLookup"
    },
    {
      "field_name": "XmlData",
      "structured_data_format_id": 1
    }
  ]
}
]
}

```

**Request (with multiple tables):**

```

{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "unload_split": 2,
        "schema_name": "DLPXDBORA",
        "table_name": "test_multi_0"
      },
      "target": {
        "stream_size": 65536,
        "schema_name": "DLPXDBORA",
        "table_name": "test_multi_0"
      },
      "masking_inventory": [
        {
          "field_name": "col_VARCHAR",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "FirstNameLookup"
        },
        {
          "field_name": "XmlData",
          "structured_data_format_id": 1
        }
      ]
    }
  ]
}

```

```

},
{
  "source": {
    "unload_split": 2,
    "schema_name": "DLPXDBORA",
    "table_name": "test_multi_1"
  },
  "target": {
    "stream_size": 65536,
    "schema_name": "DLPXDBORA",
    "table_name": "test_multi_1"
  },
  "masking_inventory": [
    {
      "field_name": "COL_TIMESTAMP",
      "domain_name": "DOB",
      "algorithm_name": "DateShiftVariable",
      "date_format": "yyyy-MM-dd HH:mm:ss.SSS" -->(optional field, this needs to be added
      only while working with date/time masking)
    }
  ]
}
]
}

```

**Response (with multiple tables):**

```

{
  "id": 1,
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "unload_split": 2,
        "schema_name": "DLPXDBORA",
        "table_name": "test_multi_0"
      },
      "target": {
        "stream_size": 65536,
        "schema_name": "DLPXDBORA",
        "table_name": "test_multi_0"
      },
      "masking_inventory": [
        {
          "field_name": "col_VARCHAR",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "FirstNameLookup"
        },
        {
          "field_name": "XmlData",
          "structured_data_format_id": 1
        }
      ]
    }
  ]
}

```

```

}
]
},
{
"source": {
"unload_split": 2,
"schema_name": "DLPXDBORA",
"table_name": "test_multi_1"
},
"target": {
"stream_size": 65536,
"schema_name": "DLPXDBORA",
"table_name": "test_multi_1"
},
"masking_inventory": [
{
"field_name": "COL_TIMESTAMP",
"domain_name": "DOB",
"algorithm_name": "DateShiftVariable",
"date_format": "yyyy-MM-dd HH:mm:ss.SSS"
}
]
}
]
}
}

```

**Delimited Files DataSet request:**

```

{
"connector_id": 1,
"mount_filesystem_id": 1,
"data_info": [
{
"source": {
"delimiter": "|",
"endOfRecord": "\n",
"enclosure": "\"",
"enclosureEscapeCharacter": "\\",
"escapeEnclosureEscapeCharacter": false,
"unique_source_files_identifier": "file_identifier1",
"has_headers": false,
"unload_split": 100,
"source_files": [
"file1.txt",
"file2.txt"
]
},
"target": {
"perform_join": true
},
"masking_inventory": [
{

```

```

    "field_name": "f3",
    "domain_name": "FIRST_NAME",
    "algorithm_name": "FirstNameLookup"
  },
  {
    "field_name": "f5",
    "domain_name": "LAST_NAME",
    "algorithm_name": "LastNameLookup"
  }
]
}

```

**The source DataSet consists of the following parameters:**

1. **delimiter:** The single character length delimiter used in source files.
2. **endOfRecord :** The end of line character, currently we only support ``\n``, ``\r`` & ``\r\n``.
3. **enclosure:** The single character length quote character used in the source files.
4. **enclosureEscapeCharacter:** The escape character used to escape quote characters.
5. **unique\_source\_files\_identifier:** This is the source key that maps the load-service and masking-service data sets with the unload-service data set. Please ensure that this value is different for each item in the DataSet data\_info list.
6. **has\_headers:** A flag that indicates if the character source files have header column names or not.
  - a. If set to **true**, format files with the same column names are created and the same can be used for the masking inventory.
  - b. If set to **false**, the column names of pattern **f0**, **f1**, **f2**, and so on are used to create the format files for delimited file masking. When adding the masking inventory please make sure to use **field\_name** with values **f0**, **f1**, **f2**, and so on.
7. **unload\_split:** The number of splits that the files in the source\_files list have to be split into. Please ensure that the split number is not too small nor too big for better overall performance.
8. **source\_files:** List of all source files that need to be masked and adhere to the following rules:
  - a. All files should have the same delimiter character and other helper characters.
  - b. All files should have the same number of columns and the same column names if it has a header line.
  - c. Supported input formats:
    - i. Relative path to the file - relative to the configured connector source path. Examples (considering that the source path within the container is `/mnt/source`):
      1. If the absolute path of the file was ``/mnt/source/file1.txt``, then the source\_files list should only contain ``file1.txt``.
      2. If the absolute path of the file was ``/mnt/source/some_dir/file1.txt``, then the source\_files list should contain the value ``some_dir/file1.txt``.
    - ii. Blob pattern - the path to all files matching a blob. Examples:



1. If we want to mask files ``/mnt/source/file1.txt`` and ``/mnt/source/file2.txt``, then the `source_files` list can contain the value ``file*.txt``.
2. If we want to mask all files within a directory, we can add ``directory/*`` to the `source_files` list.

**The target DataSet consists of the following parameters:**

1. `perform_join`: A flag to check if the load-service joins back all masked files or not.
  - a. If set to true, the load-service will join back masked files and keep the same relative file location structure in the target path (considering that the target path within the container is `/mnt/target`).
    - i. If the input was ``file1.txt``, then the output file will be ``/mnt/target/file1.txt``.
    - ii. If the input was ``some_dir/file1.txt``, then the output file will be ``/mnt/target/some_dir/file1.txt``.
  - b. If set to false, the split masked files will be placed in the same relative file location structure in the target path with a split number appended to the file name.
    - i. If the input was ``file1.txt``, then the output files will be ``/mnt/target/file1_0.txt``, ``/mnt/target/file1_1.txt``, and so on.
    - ii. If the input was ``some_dir/file1.txt``, then the output files will be ``/mnt/target/some_dir/file1_0.txt``, ``/mnt/target/some_dir/file1_1.txt``, and so on.

The masking inventory remains the same, except when there are no column names in the delimited files and the connector assumes that the column names would be ``f0``, ``f1``, ``f2``, and so on. These are **field names** that have to be used to fill up the masking inventory (as shown in the example above).

**Delimited Files DataSet response:**

```
{
  "id": 1,
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "delimiter": "|",
        "endOfRecord": "\n",
        "enclosure": "\"",
        "enclosureEscapeCharacter": "\\\"",
        "escapeEnclosureEscapeCharacter": false,
        "unique_source_files_identifier": "file_identifier1",
        "has_headers": false,
        "unload_split": 100,
        "source_files": [
          "file1.txt",
```

```

        "file2.txt"
    ]
},
"target": {
    "perform_join": true
},
"masking_inventory": [
    {
        "field_name": "f3",
        "domain_name": "FIRST_NAME",
        "algorithm_name": "FirstNameLookup"
    },
    {
        "field_name": "f5",
        "domain_name": "LAST_NAME",
        "algorithm_name": "LastNameLookup"
    }
]
}
]
}

```

**MongoDB DataSet request:**

```

{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "database_name": "SRC_DB",
        "collection_name": "SRC_COLLECTION",
        "unload_split": 10
      },
      "target": {
        "database_name": "TGT_DB",
        "collection_name": "TGT_COLLECTION"
      },
      "masking_inventory": [
        {
          "field_name": "$[*]['relationships'][*]['person']['first_name']",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "dlpx-core:FullName"
        },
        {
          "field_name": "$[*]['address']",
          "domain_name": "ADDRESS",
          "algorithm_name": "dlpx-core:CM Alpha-Numeric"
        }
      ]
    }
  ]
}
]

```

```
}

```

- **unload\_split**: The number of splits that the MongoDB collection has to be split into. You must make sure that the split number is neither too small nor too big for a better overall performance.
- **masking\_inventory**: The masking inventory for MongoDB collection can be generated by using the MongoDB Profiler service, which will call the dataset API with the generated masking\_inventory and create the dataset. MongoDB Profiler artifact includes a README file that provides detailed usage instructions.

#### MongoDB DataSet response:

```
{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "database_name": "SRC_DB",
        "collection_name": "SRC_COLLECTION",
        "unload_split": 10
      },
      "target": {
        "database_name": "TGT_DB",
        "collection_name": "TGT_COLLECTION"
      },
      "masking_inventory": [
        {
          "field_name": "$[*]['relationships'][*]['person']['first_name']",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "dlpx-core:FullName"
        },
        {
          "field_name": "$[*]['address']",
          "domain_name": "ADDRESS",
          "algorithm_name": "dlpx-core:CM Alpha-Numeric"
        }
      ]
    }
  ]
}
```

#### Parquet DataSet request:

```
{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "unique_source_files_identifier": "file_identifier1",

```

```

"unload_split": 100,
"file_type": "parquet",
"source_files": [
  "relative_folder1/file1.parquet",
  "relative_folder2/file2.parquet"
]
},
"target": {
  "perform_join": true
},
"masking_inventory": [
  {
    "field_name": "column_name1",
    "domain_name": "FIRST_NAME",
    "algorithm_name": "FirstNameLookup"
  },
  {
    "field_name": "column_name2",
    "domain_name": "LAST_NAME",
    "algorithm_name": "LastNameLookup"
  }
]
}
]
}

```

**i** The **Hyperscale Parquet Profiler** can be used to analyze the source parquet files and help generate the DataSet with the masking inventory. To know more, visit the [Parquet Profiler documentation](#) (see page 203).

**The source DataSet consists of the following parameters:**

1. **unique\_source\_files\_identifier:** This is the source key that maps the load-service and masking-service data sets with the unload-service data set. Please ensure that this value is different for each item in the DataSet data\_info list.
2. **unload\_split:** The number of splits that the files in the source\_files list have to be split into. Please ensure that the split number is not too small nor too big for better overall performance.
3. **file\_type:** The file type should be set to “parquet”.
4. **source\_files:** List of all source files that need to be masked and adhere to the following rules:
  - a. All files should have the same delimiter character and other helper characters.
  - b. All files should have the same number of columns and the same column names if it has a header line.
  - c. Supported input formats:
    - i. Relative path to the file, relative to the configured connector source path. Examples (considering that the source path within the container is s3\_bucket/source):

1. If the file's absolute path was ``s3_bucket/source/file1.parquet``, then the `source_files` list should only contain ``files1.parquet``.
  2. If the absolute path of the file was ``s3_bucket/source/some_dir/file1.parquet``, then the `source_files` list should contain the value ``some_dir/file1.parquet``.
- ii. Blob pattern, the path to all files matching a blob. Examples:
1. If we want to mask files ``s3_bucket/source/file1.parquet`` and ``s3_bucket/source/file2.parquet``, then the `source_files` list can contain the value ``file*.parquet``.
  2. If we want to mask all files within a directory, we can add ``directory/*`` to the `source_files` list.

**The target DataSet consists of the following parameters:**

1. **perform\_join**: A flag to check if the load-service joins back all masked files or not.
  - a. If set to true, the load-service will join back masked files and keep the same relative file location structure in the target path (considering that the target path within the container is `s3_bucket/target`).
    - i. If the input was ``file1.parquet``, then the output file will be ``s3_bucket/target/file1.parquet``
    - ii. If the input was ``some_dir/file1.parquet``, then the output file will be ``s3_bucket/target/some_dir/file1.parquet``.
  - b. If set to false, the split masked files will be placed in the same relative file location structure in the target path with a split number appended to the file name.
    - i. If the input was ``file1.parquet``, then the output files will be ``s3_bucket/target/file1_0.parquet``, ``s3_bucket/target/file1_1.parquet``, and so on.
    - ii. If the input was ``some_dir/file1.txt``, then the output files will be ``s3_bucket/target/some_dir/file1_0.parquet``, ``s3_bucket/target/some_dir/file1_1.parquet``, and so on.

**Parquet DataSet response:**

```
{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "unique_source_files_identifier": "file_identifier1",
        "unload_split": 100,
        "file_type": "parquet",
        "source_files": [
          "<relative_path>/file1.parquet",

```

```

    "<relative_path>/file2.parquet"
  ]
},
"target": {
  "perform_join": true
},
"masking_inventory": [
  {
    "field_name": "column_name1",
    "domain_name": "FIRST_NAME",
    "algorithm_name": "FirstNameLookup"
  },
  {
    "field_name": "column_name2",
    "domain_name": "LAST_NAME",
    "algorithm_name": "LastNameLookup"
  }
]
}
]
}
}

```

**i** Algorithm and Domain names to be provided in the Data Set request should be used from Continuous Compliance Engine. The Continuous Compliance Engine APIs that could be used to get these names are:

1. Get all algorithms ( GET /algorithms ) for Algorithm Names. Sample Endpoint: [https://maskingdocs.delphix.com/maskingApiEndpoints/5\\_1\\_15\\_maskingApiEndpoints.html#getAllAlgorithms](https://maskingdocs.delphix.com/maskingApiEndpoints/5_1_15_maskingApiEndpoints.html#getAllAlgorithms)
2. Get all domains ( GET /domains ) for Domain Names. Sample Endpoint: [https://maskingdocs.delphix.com/maskingApiEndpoints/5\\_1\\_15\\_maskingApiEndpoints.html#getAllDomains](https://maskingdocs.delphix.com/maskingApiEndpoints/5_1_15_maskingApiEndpoints.html#getAllDomains)

To check about extra parameters that need to be provided in the Data Set request for Date and Multi Column Algorithms, refer to Model `DataSet_masking_inventory` on the Hyperscale Compliance API Documentation page available in the API Reference section of this Documentation.

Alternatively, you can create/update the dataset using payload in a file with below end-points:

1. `POST /data-sets/file-upload`
2. `PUT /data-sets/file-upload/{dataSetId}`

Above endpoints requires a `file_upload_ref`, which can be generated via `POST /file-upload` endpoint. See <Link to API Doc>.

## 6.8 Jobs API


### POST /jobs (Create a Hyperscale Compliance job)

```
{
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "source_configs": {
    "max_concurrent_source_connection": 30
  },
  "target_configs": {
    "max_concurrent_target_connection": 30,
    "parallelism_degree": 15
  },
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}
```



- For more information on `retain_execution_data` flag, see [Cleaning Up Execution Data](#) (see page 202).
- In the case of Oracle, set **`parallelism_degree`** in the **`target_config`** to use the degree of parallelism while re-creating the indexes in the post-load step.
  - Set `{"parallelism_degree": 0}`**: Use unmodified DDL provided by Oracle.
  - Set `{"parallelism_degree": -1}`**: Remove any parallel or nonparallel clause from the DDL.
  - Set `{"parallelism_degree": any positive value}`**: Remove existing parallel degree or nonparallel clause and add `Parallel <parallelism_degree>` to the DDL.

**Response:**

 Below properties are only applicable to Oracle Datasource.

```
{
  "id": 1,
  "name": "Test_Job",
  "masking_engine_ids": [
    1,
    2,
    3
  ],
  "data_set_id": 1,
  "app_name_prefix": "Test_App",
  "env_name_prefix": "Test_Env",
  "retain_execution_data": "NO",
  "source_configs": {
    "max_concurrent_source_connection": 30
  },
  "target_configs": {
    "max_concurrent_target_connection": 30
  },
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 1024,
    "description": "Job created by Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}
```

#### Delimited Files Job request:

The Delimited Files job does not require the **source\_configs** and **target\_configs** objects.

```
{
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
  }
}
```



```

    "num_input_streams": 1
  }
}

```

**Delimited Files Job response:**

```

{
  "id": 1,
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}

```

**MongoDB Job request:**

The MongoDB job does not require the **source\_configs** and **target\_configs** objects.

```

{
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by MongoDB Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}

```

**MongoDB Job response:**

```

{
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by MongoDB Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}

```

**Parquet Job request:**

The Delimited Files job does not require the `source_configs` and `target_configs` objects.

```

{
  "name": "job_1",
  "masking_engine_ids": [
    1
  ],
  "data_set_id": 1,
  "app_name_prefix": "app",
  "env_name_prefix": "env",
  "retain_execution_data": "NO",
  "masking_job_config": {
    "max_memory": 1024,
    "min_memory": 0,
    "description": "Job created by Hyperscale Masking",
    "feedback_size": 100000,
    "stream_row_limit": 10000,
    "num_input_streams": 1
  }
}

```

**Parquet Job response:**

```

{
  "id": 1,
  "name": "job_1",
  "masking_engine_ids": [
    1
  ]
}

```

```

],
"data_set_id": 1,
"app_name_prefix": "app",
"env_name_prefix": "env",
"retain_execution_data": "NO",
"masking_job_config": {
  "max_memory": 1024,
  "min_memory": 0,
  "description": "Job created by Hyperscale Masking",
  "feedback_size": 100000,
  "stream_row_limit": 10000,
  "num_input_streams": 1
}
}

```

## 6.9 JobExecution API

### POST /executions (Create an execution of a Hyperscale job)

#### Request:

```

{
  "job_id": 1
}

```

Response: (Immediate response will be like below. Realtime response can be fetched using GET /executions/{execution\_id} endpoint)

```

{
  "id": 124,
  "job_id": 38,
  "status": "RUNNING",
  "create_time": "2023-05-04T12:43:03.444964",
  "tasks": [
    {
      "name": "Unload"
    },
    {
      "name": "Masking"
    },
    {
      "name": "Load"
    },
    {
      "name": "Post Load"
    }
  ]
}

```

#### GET /executions/{id}/summary (Returns the job execution by execution id in summarized format)

```

{
  "id": 72,
  "job_id": 5,
  "status": "SUCCEEDED",
  "create_time": "2022-12-18T13:38:43.722917",
  "end_time": "2022-12-18T13:43:16.554603",
  "total_objects": 4,
  "total_rows": 16,
  "tasks": [
    {
      "name": "Unload",
      "status": "SUCCEEDED",
      "start_time": "2022-12-18T13:38:44.184296",
      "end_time": "2022-12-18T13:38:54.972883",
      "received_objects": 4,
      "succeeded_objects": 4,
      "failed_objects": 0,
      "processing_objects": 0,
      "processed_rows": 16,
      "total_rows": 16
    },
    {
      "name": "Masking",
      "status": "SUCCEEDED",
      "start_time": "2022-12-18T13:38:51.979725",
      "end_time": "2022-12-18T13:42:58.569202",
      "received_objects": 4,
      "succeeded_objects": 4,
      "failed_objects": 0,
      "processing_objects": 0,
      "processed_rows": 16,
      "total_rows": 16
    },
    {
      "name": "Load",
      "status": "SUCCEEDED",
      "start_time": "2022-12-18T13:40:39.350857",
      "end_time": "2022-12-18T13:43:12.966492",
      "received_objects": 4,
      "succeeded_objects": 4,
      "failed_objects": 0,
      "processing_objects": 0,
      "processed_rows": 16,
      "total_rows": 16
    },
    {
      "name": "Post Load",
      "status": "SUCCEEDED",
      "start_time": "2022-12-18T13:43:12.981490",
      "end_time": "2022-12-18T13:43:15.764366",
      "metadata": [
        {


```

```

    "type": "Constraints",
    "total": 20,
    "processed": 20,
    "status": "SUCCESS",
    "start_time": "2022-12-18T13:43:12.981490",
    "end_time": "2022-12-18T13:43:15.764366"
  },
  {
    "type": "Indexes",
    "total": 10,
    "processed": 10,
    "status": "SUCCESS",
    "start_time": "2022-12-18T13:43:12.981490",
    "end_time": "2022-12-18T13:43:15.764366"
  },
  {
    "type": "Triggers",
    "total": 5,
    "processed": 5,
    "status": "SUCCESS",
    "start_time": "2022-12-18T13:43:12.981490",
    "end_time": "2022-12-18T13:43:15.764366"
  }
]
}
]
}

```

**GET /executions/{execution\_id} (Returns the job execution by execution\_id in the detailed format)**

 The execution response may initially return an approximate number of rows at the start of execution and provide actual values later during the execution.

#### Request:

```
id: 1
```

#### Response:

```

{
  "id": 1,
  "job_id": 1,
  "status": "SUCCEEDED",
  "create_time": "2023-04-26T12:34:38.012768",
  "end_time": "2023-04-26T12:37:32.410297",
  "total_objects": 1,
  "total_rows": 499999,

```

```
"tasks": [  
  {  
    "name": "Unload",  
    "status": "SUCCEEDED",  
    "start_time": "2023-04-26T12:34:38.027224",  
    "end_time": "2023-04-26T12:34:42.435849",  
    "metadata": [  
      {  
        "source_key": "dbo.test_TEMP",  
        "total_rows": 499999,  
        "status": "SUCCEEDED",  
        "unloaded_rows": 499999  
      }  
    ]  
  },  
  {  
    "name": "Masking",  
    "status": "SUCCEEDED",  
    "start_time": "2023-04-26T12:34:40.420073",  
    "end_time": "2023-04-26T12:35:12.423744",  
    "metadata": [  
      {  
        "source_key": "dbo.test_TEMP",  
        "total_rows": 499999,  
        "status": "SUCCEEDED",  
        "masked_rows": 499999  
      }  
    ]  
  },  
  {  
    "name": "Load",  
    "status": "SUCCEEDED",  
    "start_time": "2023-04-26T12:37:08.482240",  
    "end_time": "2023-04-26T12:37:22.417561",  
    "metadata": [  
      {  
        "source_key": "dbo.test_TEMP",  
        "total_rows": 499999,  
        "status": "SUCCEEDED",  
        "loaded_rows": 499999  
      }  
    ]  
  },  
  {  
    "name": "Post Load",  
    "status": "SUCCEEDED",  
    "start_time": "2023-04-26T12:37:22.426813",  
    "end_time": "2023-04-26T12:37:22.814583",  
    "metadata": [  
      {  
        "status": "SUCCEEDED",  
        "table_set": [  
          "test_TEMP_Result"  
        ]  
      }  
    ]  
  }  
]
```

```

],
"object_details": [
  {
    "type": "Triggers",
    "total": 2,
    "processed": 2,
    "status": "SUCCEEDED",
    "start_time": "2023-04-26T12:35:10.325948",
    "end_time": "2023-04-26T12:37:22.804792"
  },
  {
    "type": "Indexes",
    "total": 4,
    "processed": 4,
    "status": "SUCCEEDED",
    "start_time": "2023-04-26T12:35:10.325948",
    "end_time": "2023-04-26T12:37:22.804792"
  },
  {
    "type": "Constraints",
    "total": 5,
    "processed": 5,
    "status": "SUCCEEDED",
    "start_time": "2023-04-26T12:35:10.325948",
    "end_time": "2023-04-26T12:37:22.804792"
  }
]
}
]
}
]
}
}

```

- Only in case of execution failure, the below API can be used to restart the execution: `PUT /executions/{execution_id}/restart` (Restart a failed execution).
- The below API can be used only for manually cleaning up the execution: `DELETE /executions/{execution_id}` (Clean up the execution).

## 7 How to Sync a Hyperscale Job

### 7.1 How to import a Job from Continuous Compliance Engine

The `POST /import` endpoint is useful when you have a database masking job setup on a Continuous Compliance Engine and need to use the same masking inventory in a Hyperscale job. You can export the masking job details from a Continuous Compliance Engine and import them into the Hyperscale Compliance Orchestrator using the below steps.

1. Export the masking job from the Delphix Continuous Compliance Engine that needs to be imported on the Hyperscale Engine for the dataset preparation. For more information about exporting a job, refer to [Export the job](#)<sup>61</sup>.
2. After the job is exported, you can make a request on the Hyperscale Engine with the new `/import` API endpoint to upload the response blob along with `mount_filesystem_id` (Required) and `data_info_settings` (Optional) for the source and target dataset. This `data_info_settings` will be applicable to all the `data_info` objects in the dataset. For more information, refer to the [/import](#)<sup>62</sup> API page.

The following is an example of the `request blob`.

```
{
  "exportResponseMetadata": {
    "exportHost": "1.1.1.1",
    "exportDate": "Tue Sep 13 12:55:31 UTC 2022",
    "requestedObjectList": [
      {
        "objectIdentifier": {
          "id": 3
        },
        "objectType": "MASKING_JOB",
        "revisionHash": "2873bd283bd"
      }
    ],
    "exportedObjectList": [
      {
        "objectIdentifier": {
          "id": 2
        }
      }
    ]
  }
}
```

<sup>61</sup> [https://maskingdocs.delphix.com/Managing\\_Multiple\\_Engines\\_for\\_Masking/Example\\_User\\_Workflow/#1-export-the-job](https://maskingdocs.delphix.com/Managing_Multiple_Engines_for_Masking/Example_User_Workflow/#1-export-the-job)

<sup>62</sup> [https://hyperscale-compliance.delphix.com/4.0.0/hyperscaleComplianceApiEndpoints/4\\_0\\_0\\_hyperscaleComplianceApiEndpoints.html](https://hyperscale-compliance.delphix.com/4.0.0/hyperscaleComplianceApiEndpoints/4_0_0_hyperscaleComplianceApiEndpoints.html)



```

    "objectType": "SOURCE_DATABASE_CONNECTOR",
    "revisionHash": "8723bd8273b"
  },
  {
    "objectIdentifier": {
      "id": 4
    },
    "objectType": "DATABASE_CONNECTOR",
    "revisionHash": "273db2738vd"
  },
  {
    "objectIdentifier": {
      "id": 4
    },
    "objectType": "DATABASE_RULESET",
    "revisionHash": "f8c0997c804c"
  }
]
},
"blob": "983nd0239nd923ndf023nfd2p3nd923dn239dn293fn293fnb2",
"signature": "923nd023nd02",
"publicKey": "f203fn23fn203[fn230[f",
"mount_filesystem_id": 1,
"data_info_settings": [
  {
    "prop_key": "unload_split",
    "prop_value": "2",
    "apply_to": "SOURCE"
  },
  {
    "prop_key": "stream_size",
    "prop_value": "65536",
    "apply_to": "TARGET"
  }
]
}

```

3. The Hyperscale Engine will then process the required data object from the sync bundle and prepare the connector and data objects that are required for the hyperscale job creation.

4. The Hyperscale Engine will provide the data object identifier that can be further used as it is (after updating the passwords of the associated connector) to create a hyperscale job or if needed, can also be updated before configuring a job. The following is an example of the response .

```

{
  "data_set_id": id
}

```

- i** After successful import, you must provide the password for connectors manually. To do so, perform the following steps:
1. Get newly created data-set using `GET /data-sets/{datasetId}` to get the newly created connector-info id.
  2. Copy the `connector-id` and call the `GET /connector-info/{connectorInfoId}` and copy the response.
  3. Use the `PUT /connector-info/{connectorInfoId}` and in the body, paste the `GET` response and add the new password field with password value in the source and target to update the connector password.
  4. If the bundle is passphrase protected, then the same needs to be provided while importing the bundle in the API header as "passphrase". For more information about how to export passphrase encrypt bundle, refer to the [Export the object](#)<sup>63</sup> section.

## 7.2 How to re-import a Job from Continuous Compliance Engine

To update the existing dataset on the Hyperscale Compliance Orchestrator with a refreshed ruleset from the Continuous Compliance Engine, use `PUT /import/{datasetId}` endpoint.

Export the masking job from the Delphix Continuous Compliance Engine having refreshed ruleset and re-import the exported bundle into Hyperscale Compliance Orchestrator by providing the existing `datasetId`.

- i** The `data_info_settings` provided in this update request will only be applicable to objects(in the export bundle) which are not present in the existing dataset on Hyperscale Orchestrator.

## 7.3 Script to automatically import/re-import a Job from Continuous Compliance Engine

Hyperscale provides a utility script to automate the steps of syncing masking jobs inventory from Continuous Compliance Engine into the connector and dataset info of Hyperscale Compliance Orchestrator. This utility script is bundled with the release tar file and can be found at `<deployment_directory>/tools/import-scripts/`.

<sup>63</sup> [https://maskingdocs.delphix.com/Managing\\_Multiple\\_Engines\\_for\\_Masking/Example\\_User\\_Workflow/#3-export-the-object](https://maskingdocs.delphix.com/Managing_Multiple_Engines_for_Masking/Example_User_Workflow/#3-export-the-object)

## 7.4 How to sync global settings from a Delphix Continuous Compliance Engine

The `POST /sync-compliance-engines` endpoint is useful when you have global objects set up on a Continuous Compliance Engine and need to use the same global-objects-like algorithms in a Hyperscale job. You can export the details of the global object from a Continuous Compliance Engine and import them into the Hyperscale Compliance Orchestrator using the below steps.

1. Export the global settings from the Delphix Continuous Compliance Engine that needs to be imported on the Hyperscale Clustered Continuous Compliance Engines. For more information about exporting global settings, refer to [Syncing all Global Objects](#)<sup>64</sup>.
2. Once the bundle is exported, you can make a request on the Hyperscale Engine with the new `/sync-compliance-engines` endpoint to upload the response blob along with a list of Hyperscale Clusters Compliance Engines. For more information, refer to the [/sync-compliance-engines](#)<sup>65</sup> API page. The following is an example of the `request blob`.

```
{
  "exportResponseMetadata": {
    "exportHost": "1.1.1.1",
    "exportDate": "Tue Sep 13 12:55:31 UTC 2022",
    "requestedObjectList": [
      {
        "objectIdentifier": {
          "id": "global"
        },
        "objectType": "GLOBAL_OBJECT",
        "revisionHash": "897weqwj76"
      }
    ],
    "exportedObjectList": [
      {
        "objectIdentifier": {
          "id": 12
        },
        "objectType": "PROFILE_EXPRESSION",
        "revisionHash": "7dc67asch8a"
      },
      {
        "objectIdentifier": {
          "id": "BIOMETRIC"
        }
      }
    ]
  }
}
```

<sup>64</sup> [https://maskingdocs.delphix.com/Managing\\_Multiple\\_Engines\\_for\\_Masking/Example\\_User\\_Workflow/#syncing-all-global-objects](https://maskingdocs.delphix.com/Managing_Multiple_Engines_for_Masking/Example_User_Workflow/#syncing-all-global-objects)

<sup>65</sup> [https://hyperscale-compliance.delphix.com/4.0.0/hyperscaleComplianceApiEndpoints/4\\_0\\_0\\_hyperscaleComplianceApiEndpoints.html](https://hyperscale-compliance.delphix.com/4.0.0/hyperscaleComplianceApiEndpoints/4_0_0_hyperscaleComplianceApiEndpoints.html)

```

    "objectType": "DOMAIN",
    "revisionHash": "7edb8ewbd8w"
  },
  {
    "objectIdentifier": {
      "algorithmName": "d\lpx-core:Email SL"
    },
    "objectType": "USER_ALGORITHM",
    "revisionHash": "87h823d23d23"
  }
]
},
"blob": "39fdn23d9834fn3948f348fbw3pd9234nf9p4hf89",
"signature": "7823hd823bd8",
"publicKey": "892d3un293dn2p39db8283",
"compliance_engine_ids": [
  1,
  2
]
}

```



1. After import, if Hyperscale Clustered Continuous Compliance Engines already have same objects with same id or properties, then those objects will be overwritten.
2. If the bundle is passphrase protected, then the same needs to be provided while importing the bundle in the header as "passphrase". For more information about how to export passphrase encrypt bundle, refer to the [Export the object](#)<sup>66</sup> section.

## 7.5 Limitations

Hyperscale Job Sync feature has the following limitations:

1. The default maximum supported size for syncing a document or request is 50 MB. You have the option to customize this by mounting a custom `nginx.conf` under the volumes of the proxy service in the `docker-compose.yaml` file and specify `client_max_body_size`<sup>67</sup> with the new value. For more information, refer to [Custom Configuration](#) (see page 94).
2. Pre and post-script import from the Continuous Compliance Engine to Hyperscale Engine is not supported.
3. Import of Kerberos and Custom JDBC drivers connector-based making job is not supported.

<sup>66</sup> [https://maskingdocs.delphix.com/Managing\\_Multiple\\_Engines\\_for\\_Masking/Example\\_User\\_Workflow/#3-export-the-object](https://maskingdocs.delphix.com/Managing_Multiple_Engines_for_Masking/Example_User_Workflow/#3-export-the-object)

<sup>67</sup> [https://nginx.org/en/docs/http/nginx\\_http\\_core\\_module.html#client\\_max\\_body\\_size](https://nginx.org/en/docs/http/nginx_http_core_module.html#client_max_body_size)

## 8 How to cancel a Hyperscale job

To cancel a Hyperscale Job while the execution is running, the `/executions/{id}/cancel` endpoint of JobExecution API can be used. Here `{id}` is the Hyperscale Execution Id that needs to be cancelled.

Hyperscale Job cancellation is an async process. The progress of Job cancellation can be tracked by checking the Execution Status using `GET /execution/{id}/summary` API endpoint. As the cancellation process starts for a Hyperscale execution, the Execution Status becomes `CANCEL_INITIATED`. Whenever the cancellation process completes for the Execution, the Execution Status gets updated to `CANCELLED`.

During cancellation, the Hyperscale application:

1. Stops the processes running on the Hyperscale Orchestrator with respect to that Execution.
2. Closes the database connections made with the target database.
3. Cancels the Masking Jobs running on the Continuous Compliance Engines and waits for their cancellation to complete before marking Hyperscale Execution as `CANCELLED`.



Hyperscale doesn't stop any running processes on the target database. Hyperscale Job cancellation might leave the target database in an inconsistent state.

## 9 Configuration settings



1. Possible values of Configuration Settings having Type “Log Level” are TRACE, DEBUG, INFO, WARN, ERROR, FATAL, or OFF.
2. Commonly used properties can be configured in the `.env` file. The other properties must be configured in the `docker-compose.yaml` under the respective service environment.
3. If you define property values in `.env` and `docker-compose` file both, then values from `docker-compose` will take precedence.
4. The dataset date format should be the same as the environment variable date format value.
5. In one dataset, all the inventories should use the same date format for all date formats.

The following table lists the Hyperscale Compliance properties with their default values.

### 9.1 Commonly used properties

Group	Property name	Type	Description	Default value
Controller Service	API_KEY_CREATE	Boolean	This property is by default uncommented to have the container create a new API key and print it in the logs when starting. Since the value is in the logs, this API key should only be used to bootstrap the creation of other - more secure - API keys and be discarded. Comment it once the bootstrap key is available.	true

Group	Property name	Type	Description	Default value
	LOG_LEVEL_CONTROLLER_SERVICE	Log Level	Hyperscale logging level. This configuration controls the logging level of Hyperscale specific packages. This log level can be increased if Hyperscale service specific actions needs to be monitored closely. <b>NOTE:</b> It is recommended to keep this log level to INFO. Increasing the log level can impact application's performance.	INFO
	API_VERSION_COMPATIBILITY_STRICT_CHECK	Boolean	These properties are used to check the version compatibility. Setting this as true will enable strict comparison of API versions of different services. In strict comparison, the complete version i.e x.y.z is compared while in other case when this property is set to false, only major version(x out of x.y.z) of APIs will be compared.	false
	EXECUTION_STATUS_POLL_DURATION	Milli-seconds	Time duration in which execution status is collected from different services	120000
	LOAD_SERVICE_REQUIRE_POST_LOAD	Boolean	Set if the Post Load step needs to be executed.	true
	SKIP_UNLOAD_SPLIT_COUNT_VALIDATION	Boolean	Skip 'split count' and 'number of unload files generated' validation, while determining execution status <b>Note:</b> For connector-specific values, see <a href="#">Data Source Support</a> <sup>68</sup> .	false

<sup>68</sup> <http://delphixdocs.atlassian.net/hyperscale-compliance-9-0-0/docs/data-source-support-3>

Group	Property name	Type	Description	Default value
	SKIP_LOAD_SPLIT_COUNT_VALIDATION	Boolean	Skip 'split count' and 'number of masked files loaded by loaded' validation, while determining execution status <b>Note:</b> For connector-specific values, see <a href="#">Data Source Support</a> <sup>69</sup> .	false
	CANCEL_STATUS_POLL_DURATION	Milli-seconds	Time duration in which execution status is collected from different services for the CANCEL_INITIATED executions.	60000
	VALIDATE_UNLOADED_ROW_COUNT_FOR_STATUS	Boolean	Flag to control if row counts should be considered as a factor to decide the completion status of the unload step of an object.	true
	VALIDATE_MASKED_ROW_COUNT_FOR_STATUS	Boolean	Flag to control if row counts should be considered as a factor to decide the completion status of the masking step of an object.	true
	VALIDATE_LOADED_ROW_COUNT_FOR_STATUS	Boolean	Flag to control if row counts should be considered as a factor to decide the completion status of the load step of an object.	true
	DISPLAY_BYTES_INFO_IN_STATUS	Boolean	Flag to control the display of number of bytes in Execution Status response	false
	DISPLAY_ROW_COUNT_IN_STATUS	Boolean	Flag to control the display of number of rows in Execution Status response	true

<sup>69</sup> <http://delphixdocs.atlassian.net/hyperscale-compliance-9-0-0/docs/data-source-support-3>



Group	Property name	Type	Description	Default value
Unload Service	LOG_LEVEL_UNLOAD_SERVICE	Log Level	Hyperscale logging level. This configuration controls the logging level of Hyperscale specific packages. This log level can be increased if Hyperscale service specific actions needs to be monitored closely. NOTE: It is recommended to keep this log level to INFO. Increasing the log level can impact application's performance.	INFO
	UNLOAD_FETCH_ROWS	Number	Number of rows to be fetched from the database at a time.	10000
	CONCURRENT_EXPORT_LIMIT	Number	Number of concurrent mongoexport processes to be spawned.	10
Masking Service	LOG_LEVEL_MASKING_SERVICE	Log Level	Hyperscale logging level. This configuration controls the logging level of Hyperscale specific packages. This log level can be increased if Hyperscale service specific actions needs to be monitored closely. NOTE: It is recommended to keep this log level to INFO. Increasing the log level can impact application's performance.	INFO
	INTELLIGENT_LOADBALANCE_ENABLED	Boolean	Set this to false if need to enable round robin load balancing in place of intelligent load balancing.	true

Group	Property name	Type	Description	Default value
Load Service	LOG_LEVEL_LO AD_SERVICE	Log Level	Hyperscale logging level. This configuration controls the logging level of Hyperscale specific packages. This log level can be increased if Hyperscale service specific actions need to be monitored closely. NOTE: It is recommended to keep this log level to INFO. Increasing the log level can impact application's performance.	INFO
	SQLLDR_BLOB_ CLOB_CHAR_LEN GTH	Number	SQLLDR properties	20000

## 9.2 Other properties

Group	Property name	Type	Description	Default value
Controller Service	SOURCE_KEY_F IELD_NAMES	String	Dataset configuration. These fields/ columns are used to uniquely identify source data.	schema_name, table_name
	LOGGING_LEVEL_ROOT	Log Level	Logging configuration. This spring boot configuration controls the logging of all the packages/libraries getting used in application. NOTE: Increasing this Log Level will produce too many logs. It is recommended to keep this log level to WARN or below.	WARN
	LOGGING_FILE_NAME <sup>1</sup>	String	Log file location & name	/opt/ delphix/ logs/ hyperscale- controller.log

Group	Property name	Type	Description	Default value
	LOGGING_PATTERN_FILE	String	Logging pattern for file	<pre>%d{dd-MM-yyyy HH:mm:ss.SSS} \ [%thread \] %-5level %logger{ 36}.%M - %msg%n</pre>
	LOGGING_PATTERN_CONSOLE	String	Logging pattern for console	<pre>%d{dd-MM-yyyy HH:mm:ss.SSS} \ [%thread \] %-5level %logger{ 36}.%M - %msg%n</pre>

Group	Property name	Type	Description	Default value
	LOGGING_PATTERN_ROLLINGFILENAME <sup>1</sup>	String	Archived file location & name	/opt/delphix/logs/archived/hyperscale-controller-%d{yyyy-MM-dd}.%i.log
	LOGGING_FILE_MAXSIZE	File Size (String)	Max individual file size	5MB
	LOGGING_FILE_MAXHISTORY	Number of Days	History in days (i.e. keep 15 days' worth of history capped at 5GB total size)	15
	LOGGING_FILE_TOTALCAPSIZE	File Size (String)	Max limit the combined size of log archives	5GB
	LOGGING_LEVEL_ORG_SPRINGFRAMEWORK_WEB_FILTER_COMMON_SREQUESTLOGGINGFILTER	Log Level	This configuration controls the logging information of the HTTP requests received by Hyperscale. This is by default set to DEBUG level for logging request URIs of the Incoming Requests.	DEBUG
	API_VERSION_COMPATIBILITY_RETRY_COUNT	Number	These properties are used to check the version compatibility. Number of times to retry the comparison if the services are not compatible.	3

Group	Property name	Type	Description	Default value
	API_VERSION_COMPATIBILITY_RETRY_WAIT_TIME	Time in milliseconds	These properties are used to check the version compatibility. Time to wait before next retry if the services are not compatible.	10000
Unload Service	LOGGING_LEVEL_ROOT	Log Level	Logging configuration. This spring boot configuration controls the logging of all the packages/libraries getting used in application. NOTE: Increasing this Log Level will produce too many logs. It is recommended to keep this log level to WARN or below.	WARN
	LOGGING_FILE_NAME <sup>1</sup>	String	Log file location & name	/opt/delphix/logs/hyperscale-unload.log
	LOGGING_PATTERN_FILE	String	Logging pattern for file	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread \] %-5level %logger{36}.%M - %msg%n

Group	Property name	Type	Description	Default value
	LOGGING_PATTERN_CONSOLE	String	Logging pattern for console	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread \] %-5level %logger{36}.%M - %msg%n
	LOGGING_PATTERN_ROLLINGFILE_NAME <sup>1</sup>	String	Archived file location & name	/opt/delphix/logs/archived/hyperscale-unload-%d{yyyy-MM-dd}.%i.log
	LOGGING_FILE_MAXSIZE	File Size in String	Max individual file size	5MB
	LOGGING_FILE_MAXHISTORY	Number of Days	History in days (i.e. keep 15 days' worth of history capped at 5GB total size)	15
	LOGGING_FILE_TOTALSIZECAP	File Size in String	Max limit the combined size of log archives	5GB

Group	Property name	Type	Description	Default value
	LOGGING_LEVEL_ORG_SPRINGFRAMEWORK_WEB_FILTER_COMMONSREQUESTLOGGINGFILTER	Log Level	This configuration controls the logging information of the HTTP requests received by Hyperscale. This is by default set to DEBUG level for logging request URIs of the Incoming Requests.	DEBUG
	SPARK.DATE_TIMESTAMP_FORMAT	String	Spark date and timestamp format for unload. This property is only applicable for MS SQL unload image	yyyy-MM-dd HH:mm:ss.SSS
	SPARK.SMALL_DATE_TIMESTAMP_FORMAT	String	Spark small date and timestamp format for unload. This property is only applicable for MS SQL unload image	yyyy-MM-dd HH:mm
	UNLOAD.CONF_DISABLE_REPLICATION	Boolean	If set to true, disable database replication for source database. This property is only applicable for MS SQL unload image	true
	UNLOAD.SPARK_DRIVER_MEMORY	Integer	Spark driver memory to be consumed for unload. This property is only applicable for MS SQL unload image	90% of available memory
	UNLOAD.SPARK_DRIVER_CORES	Integer	Spark cores to be consumed for unload. This property is only applicable for MS SQL unload image	90% of available memory

Group	Property name	Type	Description	Default value
	UNLOAD_HIKARI_MAX_LIFETIME	Integer	Value in Milliseconds, Please follow maxLifetime from <a href="#">Hikari Documentation</a> <sup>70</sup> .	1800000 (30 min)
	UNLOAD_HIKARI_KEEP_ALIVE_TIME	Integer	Value in Milliseconds, Please follow keepaliveTime from <a href="#">Hikari Documentation</a> <sup>71</sup> .	300000 (5 min)
	FILE_DELIMITER	Char	Column value delimiter character. This configuration for internal use to have data from DB tables into files. But can be updated in specific scenario  Note: Use unicode char sequence for Quotation Marks characters (i.e. \u0022 unicode of double-quote) & Non-ASCII characters are <b>NOT</b> allowed	, (Single-Quote)
	FILE_ENCLOSURE	Char	This configuration for internal use to have data from DB tables into files. But can be updated in specific scenario (i.e. XML data masking)  Note: Use unicode char sequence for Quotation Marks characters (i.e. \u0022 unicode of double-quote) & Non-ASCII characters are <b>NOT</b> allowed	“ (Double-Quote)

<sup>70</sup> <https://github.com/brettwooldridge/HikariCP>

<sup>71</sup> <https://github.com/brettwooldridge/HikariCP>



Group	Property name	Type	Description	Default value
	FILE_ESCAPE_ENCLOSURE	Char	This configuration for internal use to have data from DB tables into files. But can be updated in specific scenario (i.e. XML data masking)  Note: Use unicode char sequence for Quotation Marks characters (i.e. \u0022 unicode of double-quote) & Non-ASCII characters are <b>NOT</b> allowed	" (Double-Quote)
Masking Service	LOGGING_LEVEL_ROOT	Log Level	Logging configuration. This spring boot configuration controls the logging of all the packages/libraries getting used in application. NOTE: Increasing this Log Level will produce too many logs. It is recommended to keep this log level to WARN or below.	WARN
	LOGGING_FILE_NAME <sup>1</sup>	String	Log file location & name	/opt/delphix/logs/hyperscale.log
	LOGGING_PATTERN_FILE	String	Logging pattern for file	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread \] %-5level %logger{36}.%M - %msg%n

Group	Property name	Type	Description	Default value
	LOGGING_PATTERN_CONSOLE	String	Logging pattern for console	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread \] %-5level %logger{36}.%M - %msg%n
	LOGGING_PATTERN_ROLLINGFILE_NAME <sup>1</sup>	String	Archived file location & name	/opt/delphix/logs/archived/hyperscale-%d{yyyy-MM-dd}.%i.log
	LOGGING_FILE_MAXSIZE	File Size in String	Max individual file size	5MB
	LOGGING_FILE_MAXHISTORY	Number of Days	History in days (i.e. keep 15 days' worth of history capped at 5GB total size)	15
	LOGGING_FILE_TOTALSIZECAP	File Size in String	Max limit the combined size of log archives	5GB

Group	Property name	Type	Description	Default value
	LOGGING_LEVEL_ORG_SPRINGFRAMEWORK_WEB_FILTER_COMMON_SREQUESTLOGGINGFILTER	Log Level	This configuration controls the logging information of the HTTP requests received by Hyperscale. This is by default set to DEBUG level for logging request URIs of the Incoming Requests.	DEBUG
	MASKING_ILB_QUEUEINGFACTOR	Float	This property, Queueing Factor(QF) can be used to increase the number of jobs being assigned to a Continuous Compliance Engine by the factor mentioned. (Net jobCapacity of CCE = CCE's total jobCapacity * Queueing Factor). It can be used to increase the utilization of Continuous Compliance Engines. Increasing this property value can lead to jobs getting queued on Masking Engines. <b>NOTE:</b> This property is only applicable if INTELLIGENT_LOADBALANCE_ENABLED is set to true(default value).	1.0
	MONITOR_POLL_DURATION	Time in seconds	Used to set the frequency with which Masking Engines will be polled to fetch job Status	10s
	ENGINE_FAILURE_POLL_DURATION	Time in seconds	Incase of Connection/Handshake issues or unavailability of Masking Engines, the subjobs on the engine will be marked as failed after retrying for this duration.	60s
Load Service	SQLLDR_SUCCESS_MESSAGE	String	Message printed by sqlldr on successful loading of data.	'successfully loaded.'

Group	Property name	Type	Description	Default value
	LOGGING_LEVEL_ROOT	Log Level	Logging configuration. This spring boot configuration controls the logging of all the packages/libraries getting used in application. NOTE: Increasing this Log Level will produce too many logs. It is recommended to keep this log level to WARN or below.	WARN
	LOGGING_LEVEL_COM_DELPHIX_MASKING	Log Level	Log level for driver support. This configuration controls the logging level of the Masking Driver Support package. This Log Level can be increased when Driver Support Steps of Load Process need to be monitored closely.	INFO
	LOGGING_FILE_NAME <sup>1</sup>	String	Log file location & name	/opt/delphix/logs/hyperscale-load.log
	LOGGING_PATTERN_FILE	String	Logging pattern for file	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread\] %-5level %logger{36}.%M - %msg%n

Group	Property name	Type	Description	Default value
	LOGGING_PATTERN_CONSOLE	String	Logging pattern for console	%d{dd-MM-yyyy HH:mm:ss.SSS} \[%thread \] %-5level %logger{36}.%M - %msg%n
	LOGGING_PATTERN_ROLLINGFILE_NAME <sup>1</sup>	String	Archived file location & name	/opt/delphix/logs/archived/hyperscale-load-%d{yyyy-MM-dd}.%i.log
	LOGGING_FILE_MAXSIZE	File Size in String	Max individual file size	5MB
	LOGGING_FILE_MAXHISTORY	Number of Days	History in days (i.e. keep 15 days' worth of history capped at 5GB total size)	15
	LOGGING_FILE_TOTALSIZECAP	File Size in String	Max limit the combined size of log archives	5GB

Group	Property name	Type	Description	Default value
	LOGGING_LEVEL_ORG_SPRINGFRAMEWORK_WEB_FILTER_COMMON_SREQUESTLOGGINGFILTER	Log Level	This configuration controls the logging information of the HTTP requests received by Hyperscale. This is by default set to DEBUG level for logging request URIs of the Incoming Requests.	DEBUG
	SPARK.DATE_TIMESTAMP_FORMAT	String	Spark date and timestamp format for load. This property is only applicable for MS SQL load image	yyyy-MM-dd HH:mm:ss.SSS
	SPARK.SMALL_DATE_TIMESTAMP_FORMAT	String	Spark small date and timestamp format for load. This property is only applicable for MS SQL load image	yyyy-MM-dd HH:mm
	LOAD.SPARK_BATCH_SIZE	Integer	Spark batch size for bulk load. This property is only applicable for MS SQL load image	10000
	LOAD.SPARK_JOB_TABLE_LOCK	Boolean	Spark job table lock for bulk load. This property is only applicable for MS SQL load image	true
	LOAD.CONF_DISABLE_REPLICATION	Boolean	If set to true, disable database replication for target database. This property is only applicable for MS SQL load image	true
	LOAD.SPARK_DRIVER_MEMORY	Integer	Spark driver memory to be consumed for load. This property is only applicable for MS SQL unload image.	90% of available memory
	LOAD.SPARK_DRIVER_CORES	Integer	Spark cores to be consumed for load. This property is only applicable for MS SQL unload image.	90% of available processors

Group	Property name	Type	Description	Default value
	<code>LOAD_HIKARI_MAX_LIFE_TIME</code>	Integer	Value in Milliseconds, Please follow <code>maxLifetime</code> from <a href="#">Hikari Documentation</a> <sup>72</sup> .	1800000 (30 min)
	<code>LOAD_HIKARI_KEEP_ALIVE_TIME</code>	Integer	Value in Milliseconds, Please follow <code>keepaliveTime</code> from <a href="#">Hikari Documentation</a> <sup>73</sup> .	300000 (5 min)



`spark.date.timestamp.format` and `spark.small.date.timestamp.format` values should be the same for MS SQL load/unload services. Also if masking is applied on the date/timestamp datatype column, the applied inventory date format should be the same for MS SQL load/unload data format.



- For each service, the file path(absolute) configured for `logging.file.name` and for `logging.pattern.rolling-file-name` has to be the same. This path is a path inside the respective container.
- For each service, if the log files(configured through `logging.file.name` and `logging.pattern.rolling-file-name`) need to be accessed outside the container, respective log path has to be mounted by adding volume binding of that path in `docker-compose.yaml` for that service.

<sup>72</sup> <https://github.com/brettwooldridge/HikariCP>

<sup>73</sup> <https://github.com/brettwooldridge/HikariCP>

## 10 Hyperscale Compliance API

The Hyperscale Compliance API is organized around REST. Our API has predictable resource-oriented URLs, accepts form-encoded request bodies, returns JSON-encoded responses, and uses standard HTTP response codes, authentication, and verbs.

### REST

Hyperscale Compliance API is a RESTful API. REST stands for REpresentational State Transfer. A REST API will allow you to access and manipulate a textual representation of objects and resources using a predefined set of operations to accomplish various tasks.

### JSON

Hyperscale Compliance API uses JSON (JavaScript Object Notation) to ingest and return representations of the various objects used throughout various operations. JSON is a standard format and, as such, has many tools available to help with creating and parsing the request and response payloads, respectively. Here are some UNIX tools that can be used to parse JSON - [Parsing JSON with Unix Tools](#)<sup>74</sup>. That being said, this is only the tip of the iceberg when it comes to JSON parsing and the reader is encouraged to use their method of choice.

### API Client

The various operations and objects used to interact with APIs are defined in a specification document. This allows us to utilize various tooling to ingest that specification to generate documentation and an API Client, which can be used to generate cURL commands for all operations.

## 10.1 Accessing the Hyperscale Compliance API

For accessing the Hyperscale Compliance API, see [Accessing the Hyperscale Compliance API](#) (see page 141).

## 10.2 View the API reference

To view the API client documentation, a downloadable `.html` file is available below.

---

<sup>74</sup> <https://stackoverflow.com/questions/1955505/parsing-json-with-unix-tools>



17\_0\_0\_0\_hyperscaleCompl...

[\(see page 200\)](#)

# 11 Cleaning up execution data

As part of the Hyperscale execution run some data files and objects are created which should be cleaned on the execution completion. These files and objects are:

1. The system will create data files (unload service) and masked files (masking service) on the file server. As the data size can be large (2 times of source data) and include sensitive information, therefore, it is important to clean up this data.
2. The system will create multiple data objects like connectors, rulesets, file formats, jobs, etc on the respective Continuous Compliance Engines. Objects created by Hyperscale should be cleaned once Hyperscale execution is complete.
3. Additionally unload service, masking service, and load service will also store transient internal data for the execution while running it. This data is not required once execution is completed.

Following are the three ways this data will be/can be cleaned.

## 1. Using retain\_execution\_data

While setting up a Hyperscale Job ( `POST /jobs` ), you can set the value for `retain_execution_data` property to the intimate system when it should clean up data automatically based on the table below.

EXECUTION_STATUS	RETAIN_EXECUTION_DATA	CLEAN UP AUTOMATICALLY?
NA(SUCCESS/FAILED/CANCELED)	NO	YES
SUCCESS	ON_ERROR	YES
FAILED	ON_ERROR	NO
CANCELED	ON_ERROR	NO
NA(SUCCESS/FAILED/CANCELED)	ALWAYS	NO

## 2. Manual clean up

Hyperscale exposes a delete API ( `DELETE /executions/{id}` ) to manually clean up data for execution if it's not already cleaned.

## 3. Start a new execution

While starting a new execution, Hyperscale will first validate if the previous execution data is cleaned. If it's not cleaned, then Hyperscale will trigger cleanup before starting new execution.

## 12 Hyperscale profilers

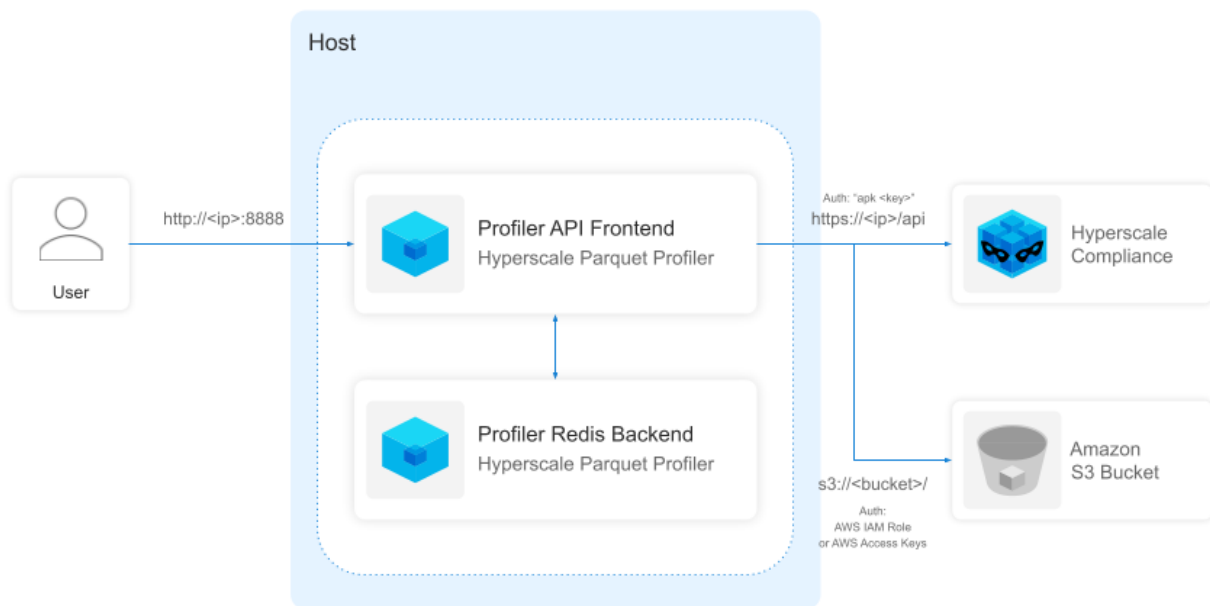
This section covers the following topics:

- [Parquet profiler \(see page 203\)](#)

### 12.1 Parquet profiler

Identifying the schema of parquet files and aligning columns with the suitable masking algorithms from the Delphix Compliance Engine poses a challenge. Moreover, when dealing with a large number of files (similar/ varied) in the source location, establishing a Hyperscale Parquet Connector dataset becomes challenging. The Parquet Profiler serves as a solution to address this complexity.

The main objective of the Parquet Profiler is to analyze the source parquet files and create a Hyperscale Parquet Connector dataset, essentially helping in creating a masking inventory for the source data.



1 Hyperscale Parquet Profiler - Deployment Architecture



The Parquet Profiler is tightly coupled with the Hyperscale Parquet Connector and cannot be used for other connectors.

## 12.1.1 Installation and setup (Parquet profiler)

### 12.1.1.1 Pre-requisites

1. Download the latest version of the Parquet Profiler from the Delphix [download](#)<sup>75</sup> page.
2. Ensure that the host running the profiler has `docker-compose` installed (the profiler has been tested only in the `docker-compose` environment).
3. Currently, the profiler supports AWS S3 buckets as the source locations. We need to ensure that the profiler has access to the source location (similar to how access was set up for the Hyperscale Parquet Connector). You can use the following authentication mechanisms:
  - Attaching the EC2 host running the profiler with an AWS IAM role which has access to the source S3 buckets.
    - IAM Roles are designed for applications to securely make AWS-API requests from EC2 instances, without the necessity to manage the security credentials that the applications use.
    - Using the AWS console UI or AWS CLI, attach the IAM role to the EC2 instance running the Hyperscale services. To know more, check the [AWS Documentation](#)<sup>76</sup>.
  - Generating an AWS Access Key ID & AWS Secret Access Key pair for an AWS Role which has the privileges to access the source S3 bucket.
    - Access keys are long-term credentials generated for an IAM user or role. These keys can be for programmatic requests to the AWS CLI or AWS API (directly or using the AWS SDK). For more information, refer to the [AWS Documentation](#)<sup>77</sup>.
4. Set up the [Hyperscale Parquet Connector](#)<sup>78</sup> and add the required [MountFileSystems](#)<sup>79</sup> and [ConnectorInfo](#)<sup>80</sup> details.

### 12.1.1.2 Procedure

1. Untar the profiler downloaded from Delphix's download page. It should contain the docker images for the profiler and the `docker-compose.yaml` file to run the profiler.

---

<sup>75</sup> <https://download.delphix.com/folder/2224/Delphix%20Product%20Releases/Hyperscale%20Compliance>

<sup>76</sup> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

<sup>77</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_access-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)

<sup>78</sup> <https://hyperscalemasking.delphix.com/docs/latest/deployment>

<sup>79</sup> [https://hyperscalemasking.delphix.com/docs/latest/how-to-setup-a-hyperscale-compliance-job#id-\(15.0.0\)HowtsetupaHyperscaleCompliancejob-MountFileSystemsAPI](https://hyperscalemasking.delphix.com/docs/latest/how-to-setup-a-hyperscale-compliance-job#id-(15.0.0)HowtsetupaHyperscaleCompliancejob-MountFileSystemsAPI)

<sup>80</sup> [https://hyperscalemasking.delphix.com/docs/latest/how-to-setup-a-hyperscale-compliance-job#id-\(15.0.0\)HowtsetupaHyperscaleCompliancejob-ConnectorInfoAPI](https://hyperscalemasking.delphix.com/docs/latest/how-to-setup-a-hyperscale-compliance-job#id-(15.0.0)HowtsetupaHyperscaleCompliancejob-ConnectorInfoAPI)

```
tar -xf parquet-profiler.tar.gz
```

2. Load the `delphix-hyperscale-profiler-api` and `delphix-hyperscale-profiler-backend` docker images.

```
docker load --input delphix-hyperscale-parquet-profiler-api.tar
docker load --input delphix-hyperscale-parquet-profiler-backend.tar
```

3. Edit the `docker-compose` YAML file to map the controller end-point for the `delphix-hyperscale-profiler-api` to interact with.

```
services:
  ...
  profiler-api-service:
  ...
  environment:
  ...
  - CONTROLLER_URL=https://<controller-ip>/api
```

4. [Optional] You can provide the AWS Access keys as environment variables as well, it will be considered as the default credentials to access the source S3 location.

```
services:
  ...
  profiler-api-service:
  ...
  environment:
  ...
  - AWS_ACCESS_KEY_ID=<access_key_id>
  - AWS_SECRET_ACCESS_KEY=<secret_access_key>
  - AWS_DEFAULT_REGION=<region>
```

5. Start the profiler service.

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

6. Access the profiler swagger UI at `http://<host-ip>:8888`.

## 12.1.2 Executing a profiler task

### 12.1.2.1 Perform the following steps to execute a profiler task

1. Add the authorization key [generated for the controller service](#)<sup>81</sup> into the profiler UI.
  - Click on the **Authorize** button and then add the key as follows, "apk <authorization-key>"
2. Validate all the configured connectors on the Hyperscale controller using the `/connector-info` GET API endpoint.
3. If the connector-info contains the AWS credentials, then the response will have the AWS credentials hidden.

**Example: /connector-info response with AWS credentials:**

```
[
  {
    "source": {
      "type": "AWS",
      "properties": {
        "server": "S3",
        "path": "s3_bucket_source/sub_folder",
        "aws_region": "us-east-1",
        "aws_access_key_id": "AKIA*****",
        "aws_secret_access_key": "x2IX*****",
        "aws_role_arn": "56436882398"
      }
    },
    "target": {
      "type": "AWS",
      "properties": {
        "server": "S3",
        "path": "s3_bucket_target/sub_folder",
        "aws_region": "us-east-1",
        "aws_access_key_id": "AKIA*****",
        "aws_secret_access_key": "x2IX*****",
        "aws_role_arn": "56436882398"
      }
    }
  }
]
```

<sup>81</sup> <https://hyperscalemasking.delphix.com/docs/latest/accessing-the-hyperscale-compliance-api>

As the credentials are masked, the profiler will need the credentials independently (in case the IAM role-based authentication is not used or the AWS credentials are not set using the environment variables).

Use the `/source-credentials/{connectorId}` post API endpoint to add the credentials mapped to the connector ID received from the controller.

#### POST /source-credentials/{connectorId} - Request JSON

```
{
  "aws_access_key_id": "AKIAJSJDFJSBSG",
  "aws_secret_access_key": "x2IXHFkDjSkdnldf&kksdfh%jsdf"
}
```

#### POST /source-credentials/{connectorId} - Response JSON

```
{
  "connector_id": 1,
  "aws_access_key_id": "AKIA*****",
  "aws_secret_access_key": "x2IX*****"
}
```

4. Validate all the mount-file systems configured in the controller using the `/mount-file systems` API endpoint.
5. The profile sets are essentially a list of all masking algorithms mapped to domain names which the profiler can assign to columns. No default profile set is created when starting the Parquet Profiler for the first time. To create a default profile set, hit the API endpoint `/profile-sets`. There should now be a default profile set with ID 1.

#### GET /profile-sets - Response JSON

```
[
  {
    "exclusions": [
      "_id",
      "_id.oid",
      "$oid",
      "_id.$oid",
      "id"
    ],
    "set_id": 1,
    "date_created": "2023-12-14T12:46:34.686136",
    "name": "DEFAULT",
    "description": "default profiler set",
    "entities": [
      {
        "domain_name": "ZIP",
        "algorithm_name": "dlpx-core:CM Alpha-Numeric",

```

```

    "type": "pattern",
    "regex": "(\\b\\d{5}(?:\\-\\d{4})?\\b)",
    "meta_context": [
      "zip",
      "code"
    ]
  },
  {
    "domain_name": "CREDIT_CARD",
    "algorithm_name": "CreditCard",
    "type": "DL"
  },
  {
    "domain_name": "DOB",
    "algorithm_name": "DateShiftDiscrete",
    "date_format": "yyyy-mm-dd",
    "type": "DL_DT",
    "min_age_years": 18,
    "max_age_years": 100
  },
  {
    "domain_name": "EMAIL",
    "algorithm_name": "dlpx-core:Email SL",
    "type": "DL"
  },
  {
    "domain_name": "IP_ADDRESS",
    "algorithm_name": "dlpx-core:CM Alpha-Numeric",
    "type": "DL"
  },
  {
    "domain_name": "ADDRESS",
    "algorithm_name": "AddrLookup",
    "type": "DL"
  },
  {
    "domain_name": "CITY",
    "algorithm_name": "USCitiesLookup",
    "type": "DL"
  },
  {
    "domain_name": "COUNTRY",
    "algorithm_name": "NullValueLookup",
    "type": "DL"
  },
  {
    "domain_name": "FIRST_NAME",
    "algorithm_name": "dlpx-core:FirstName",
    "type": "DL"
  },
  {
    "domain_name": "LAST_NAME",

```



```

    "algorithm_name": "dlpx-core:LastName",
    "type": "DL"
  },
  {
    "domain_name": "FULL_NAME",
    "algorithm_name": "dlpx-core:FullName",
    "type": "DL"
  },
  {
    "domain_name": "TELEPHONE_NO",
    "algorithm_name": "dlpx-core:Phone US",
    "type": "DL"
  },
  {
    "domain_name": "WEB",
    "algorithm_name": "WebURLsLookup",
    "type": "DL"
  },
  {
    "domain_name": "DRIVING_LC",
    "algorithm_name": "DrivingLicenseNoLookup",
    "type": "DL"
  },
  {
    "domain_name": "SSN",
    "algorithm_name": "dlpx-core:CM Alpha-Numeric",
    "type": "DL"
  }
],
"date_last_updated": "2023-12-14T12:46:34.686142"
}
]

```

6. Generally, the default profile set should be enough for most use cases. But if you want to map different masking algorithms available in your Delphix Compliance Engine to different domains, you should create your own profile set using the `/profile-sets` POST API endpoint. To know more about the profile sets available in your Delphix Compliance Engine, visit [here](#)<sup>82</sup>.

#### POST /profile-sets - Request JSON

```

{
  "set_id": 2,
  "name": "custom_profile_set",
  "description": "Different Algorithm Mapping",
  "exclusions": [
    "_id",
    "_id.oid",
    "$oid",
  ]
}

```

<sup>82</sup> <https://masking.delphix.com/docs/latest/identifying-sensitive-data>

```

    "_id.$oid",
    "id"
  ],
  "entities": [
    {
      "domain_name": "FIRST_NAME",
      "algorithm_name": "dlpx-core:FirstName",
      "type": "DL"
    },
    {
      "domain_name": "LAST_NAME",
      "algorithm_name": "dlpx-core:LastName",
      "type": "DL"
    }
  ]
}

```

Understanding the profile-set payload parameters:

- a. **name:** Name of the profile set.
- b. **exclusions:** List of fields (or column names) to exclude from the discovery.
- c. **entities:** List of entity types to run discovery:
  - i. **domain\_name:** The domain name must exist in the Compliance Engine. Note, any DL type entities Domain Name cannot be modified
  - ii. **algorithm\_name:** Any available algorithm whether out of the box or custom can be assigned to any entity type
  - iii. **type:** These are the following types of entities are allowed:
    1. **“DL”:** A Deep Learning & NLP based discovery. All DL entities must have their correspondence Domain Name from the table listed [here \(see page 0\)](#). Example payload:

```

{
  "domain_name": "CREDIT_CARD",
  "algorithm_name": "CreditCard",
  "type": "DL"
}

```

2. **“context”:** Where users can provide their list of explicit values for discovery. Example payload:

```

{
  "domain_name": "TITLE",
  "algorithm_name": "RandomValueLookup",
  "type": "context",
  "list": [
    "Mr.",
    "Mrs.",
    "Ms.",
    "Miss",
  ]
}

```

```

    "Madam",
    "Master"
  ]
}

```

3. **"pattern"** - The regex-based entity, users can add their regex criteria. Additionally, a list of fields can be supplied to provide further context to support regex discovery. Example payload:

```

{
  "domain_name": "ZIP_CODE",
  "algorithm_name": "dlpx-core:CM Alpha-Numeric",
  "type": "pattern",
  "regex": "(\\b\\d{5}(?:\\-\\d{4})?\\b)",
  "meta_context": [
    "zip",
    "code"
  ]
}

```

#### POST /profile-sets - Response JSON

```

{
  "set_id": 2,
  "name": "custom_profile_set",
  "description": "Different Algorithm Mapping",
  "exclusions": [
    "_id",
    "_id.oid",
    "$oid",
    "_id.$oid",
    "id"
  ],
  "entities": [
    {
      "domain_name": "FIRST_NAME",
      "algorithm_name": "dlpx-core:FirstName",
      "type": "DL"
    },
    {
      "domain_name": "LAST_NAME",
      "algorithm_name": "dlpx-core:LastName",
      "type": "DL"
    }
  ]
}

```

7. You can now start a profiler task using the `/tasks` POST API endpoint.

#### POST /tasks - Request JSON

```
{
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "set_id": 1,
  "scan_depth": 1000,
  "unique_source_files_identifier": "file_identifier",
  "unload_split": 2,
  "file_type": "parquet"
}
```

Understanding the task payload parameters:

- a. **connector\_id** - The connector to get the source details from. The profiler will identify all files (recursively) within the source S3 path provided in the connector-info details.
- b. **mount\_filesystem\_id** - The mount filesystem ID that the resultant Hyperscale Parquet Connector dataset should be populated with.
- c. **set\_id** - The profiler set ID that the profiler tasks should run against.
- d. **scan\_depth** - The number of (random) rows in the parquet file that need to be analyzed by the profiler to determine what kind of sensitive data it is.
- e. **unique\_source\_files\_identifier** - The source key value that the resultant Hyperscale Parquet Connector dataset should be populated with.
- f. **unload\_split** - The unload split that the resultant Hyperscale Parquet Connector dataset should be populated with.
- g. **file\_type** - The file type should be **"parquet"**.

#### POST /tasks - Response JSON

```
{
  "task_id": "11b92f0f-7c08-4768-97c5-17ce73213dc8",
  "status": "RUNNING"
}
```

8. The status of the task can be monitored using the `/tasks/{id}` GET API endpoint.
9. Once the status shows "SUCCESS", the Hyperscale Parquet Connector dataset generated by the profiler is shown as part of the results.

#### GET /tasks/{id} - Response JSON

```
{
  "task_id": "11b92f0f-7c08-4768-97c5-17ce73213dc8",
  "connector_id": 1,
  "data_set_id": null,
  "mount_filesystem_id": 1,
  "status": "SUCCESS",
  "set_id": 1,
  "scan_depth": 100,
  "file_type": "parquet",
  "unique_source_files_identifier": "file_identifier",
}
```

```

"unload_split": 2,
"results": {
  "connector_id": 1,
  "mount_filesystem_id": 1,
  "data_info": [
    {
      "source": {
        "unique_source_files_identifiser": "file_identifiser_1",
        "file_type": "parquet",
        "unload_split": 2,
        "source_files": [
          "customer/part-00000.gz.parquet",
          "customer/part-00001.gz.parquet",
          "customer/part-00002.gz.parquet",
          "customer/part-00003.gz.parquet",
          "customer/part-00004.gz.parquet",
          "customer/part-00005.gz.parquet",
          "customer/part-00006.gz.parquet",
          "customer/part-00007.gz.parquet",
          "customer/part-00008.gz.parquet",
          "customer/part-00009.gz.parquet"
        ]
      },
      "target": {
        "perform_join": true
      },
      "masking_inventory": [
        {
          "field_name": "c_last",
          "domain_name": "FIRST_NAME",
          "algorithm_name": "dlpx-core:FirstName"
        },
        {
          "field_name": "c_state",
          "domain_name": "LAST_NAME",
          "algorithm_name": "dlpx-core:LastName"
        },
        {
          "field_name": "c_phone",
          "domain_name": "TELEPHONE_NO",
          "algorithm_name": "dlpx-core:Phone US"
        }
      ]
    },
    {
      "source": {
        "unique_source_files_identifiser": "file_identifiser_2",
        "file_type": "parquet",
        "unload_split": 2,
        "source_files": [
          "district/part-00000.gz.parquet"
        ]
      }
    }
  ]
}

```

```

},
"target": {
  "perform_join": true
},
"masking_inventory": [
  {
    "field_name": "d_name",
    "domain_name": "LAST_NAME",
    "algorithm_name": "dlpx-core:LastName"
  },
  {
    "field_name": "d_street_2",
    "domain_name": "LAST_NAME",
    "algorithm_name": "dlpx-core:LastName"
  },
  {
    "field_name": "d_state",
    "domain_name": "LAST_NAME",
    "algorithm_name": "dlpx-core:LastName"
  }
]
},
{
  "source": {
    "unique_source_files_identifier": "file_identifier_7",
    "file_type": "parquet",
    "unload_split": 2,
    "source_files": [
      "orders/part-00000.gz.parquet",
      "orders/part-00001.gz.parquet",
      "orders/part-00002.gz.parquet",
      "orders/part-00003.gz.parquet",
      "orders/part-00004.gz.parquet"
    ]
  },
  "target": {
    "perform_join": true
  },
  "masking_inventory": [
    {
      "field_name": "o_id",
      "domain_name": "TELEPHONE_NO",
      "algorithm_name": "dlpx-core:Phone US"
    }
  ]
},
{
  "source": {
    "unique_source_files_identifier": "file_identifier_9",
    "file_type": "parquet",
    "unload_split": 2,
    "source_files": [

```

```

        "warehouse/part-00000.gz.parquet"
    ]
},
"target": {
    "perform_join": true
},
"masking_inventory": [
    {
        "field_name": "w_name",
        "domain_name": "CITY",
        "algorithm_name": "USCitiesLookup"
    },
    {
        "field_name": "w_street_1",
        "domain_name": "ZIP",
        "algorithm_name": "dlpx-core:CM Alpha-Numeric"
    },
    {
        "field_name": "w_state",
        "domain_name": "LAST_NAME",
        "algorithm_name": "dlpx-core:LastName"
    },
    {
        "field_name": "w_zip",
        "domain_name": "LAST_NAME",
        "algorithm_name": "dlpx-core:LastName"
    }
]
}
]
}
"total": 16,
"identified": null,
"completion": 100,
"elapsed_time": "0:06:47.837970",
"start_time": "2023-12-14T13:32:18.913943",
"end_time": "2023-12-14T13:39:06.756026",
"date_created": "2023-12-14T13:32:18.913948",
"date_last_updated": "2023-12-14T13:32:18.913950"
}

```

10. You can push the generated dataset directly from the profiler using the `/data-sets/{task_id}` POST API endpoint. The response contains the ID of the newly created dataset on the controller.

#### POST /data-sets/{task\_id} - Response JSON

```

{
  "data_set_id": 1
}

```

## 12.1.2.2 The DL entities within the default Profiler-Set with their algorithms

Type	Domain Name	Algorithm	Description
DL	FULL_NAME	dlpx-core:FullName	Full name detection
DL	FIRST_NAME	dlpx-core:FirstName	First name
DL	LAST_NAME	dlpx-core:LastName	Last name
DL	EMAIL	dlpx-core:Email SL	Email address
DL	TELEPHONE_NO	dlpx-core:Phone US	Phone or Mobile number
DL	DOB	DateShiftDiscrete	Date of Birth
DL	IP ADDRESS	dlpx-core:CM Alpha-Numeric	IP Address
DL	CREDIT CARD	CreditCard	Credit Card
DL	ADDRESS	AddrLookup	Street Address
DL	CITY	USCitiesLookup	City name
DL	COUNTRY	NullValueLookup	Country name
DL	WEB	WebURLsLookup	URL or domain name
DL	DRIVING_LC	DrivingLicenseNoLookup	US driving license
DL	SSN	dlpx-core:CM Alpha-Numeric	Social Security Number

**The other available DL entities:**

Type	Domain Name	Description
------	-------------	-------------



DL	STATE	State name
DL	STATE_CODE	State Code
DL	CRYPTO	Bitcoin address
DL	IBAN_CODE	The International Bank Account Number (IBAN)
DL	US_BANK_NUMBER	A US bank account number is between 8 to 17 digits.
DL	US_ITIN	US Individual Taxpayer Identification Number (ITIN
DL	US_PASSPORT	A US passport number with 9 digits