



Hewlett Packard
Enterprise

HPE Structured Data Manager

Software Version: 7.52

Runtime Guide

Document Release Date: December 2017

Software Release Date: December 2017

Legal notices

Warranty

The only warranties for Hewlett Packard Enterprise Development LP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted rights legend

Confidential computer software. Valid license from HPE required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright notice

© Copyright 2015-2017 Hewlett Packard Enterprise Development LP

Trademark notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

This product includes an interface of the 'zlib' general purpose compression library, which is Copyright © 1995-2002 Jean-loup Gailly and Mark Adler.

Documentation updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent software updates, go to <https://softwaresupport.hpe.com/patches>.

To verify that you are using the most recent edition of a document, go to <https://softwaresupport.hpe.com/manuals>.

This site requires that you register for an HPE Passport and sign in. To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support Online web site at <https://softwaresupport.hpe.com>.

This web site provides contact information and details about the products, services, and support that HPE Software offers.

HPE Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Access product documentation
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HPE Passport user and sign in. Many also require a support contract.

To register for an HPE Passport ID, go to <https://hpp12.passport.hpe.com/hppcf/login.do>.

To find more information about access levels, go to <https://softwaresupport.hpe.com/web/softwaresupport/access-levels>.

Contents

| | |
|---|--------|
| About this document | 10 |
| Intended audience | 10 |
| Prerequisites | 10 |
| New and revised information | 10 |
| Related documentation | 11 |
| Part I: Introduction | 13 |
| Chapter 1: Introduction to structured data management | 14 |
| Before you begin | 14 |
| Web Console overview | 15 |
| Part II: Tutorials | 16 |
| Chapter 2: Configure the Web Console | 17 |
| Start the Web Console | 17 |
| Install the repository | 19 |
| Create an environment | 24 |
| Configure settings | 33 |
| Set up email access | 33 |
| Manage indexing server settings | 34 |
| Configure the trusted store for certificates | 35 |
| Masking Server settings | 35 |
| Prerequisites | 35 |
| Date-Time Formats in SDM | 37 |
| Data Access Cartridge | 38 |
| Manage the home directory | 39 |
| Change the home directory | 39 |
| Create and edit Web Console users and groups (optional) | 40 |
| Create or edit a user group | 40 |
| Create or edit a user | 41 |
| Verify the new Web Console user | 41 |
| Chapter 3: Deploy and run business flows | 43 |
| Before you begin | 43 |
| Deploy business flows | 44 |
| Run a business flow | 46 |
| Run an undo business flow | 47 |
| Run a reload business flow | 47 |
| Before you begin | 48 |
| Chapter 4: Upload to HPE Vertica | 49 |

| | |
|--|----|
| Introduction | 49 |
| Create a Vertica location | 50 |
| Deploy the business flow | 50 |
| Specify a schema mapping file | 50 |
| Run the business flow | 51 |
| Access the archived data | 52 |
| Chapter 5: Perform multiple data movements | 53 |
| Introduction | 53 |
| Configure HPE Structured Data Manager | 53 |
| Configure a gateway instance of HPE Structured Data Manager | 54 |
| Configure an application instance of HPE Structured Data Manager | 55 |
| Run business flows | 56 |
| Chapter 6: View eligibility analytics | 58 |
| Before you begin | 58 |
| Enable eligibility analytics | 58 |
| Run a business flow | 59 |
| View eligibility analytics | 60 |
| Continue the paused business flow | 61 |
| Run the undo business flow | 61 |
| Part III: Reference guide | 62 |
| Chapter 7: Start and configure the Web Console | 63 |
| Start and stop the Web Console | 63 |
| Start the Web Console | 63 |
| Before starting the Web Console | 63 |
| Start Web Console from the program group | 63 |
| Start Web Console from the command line | 64 |
| Start Web Console as a Windows service | 64 |
| Start Web Console on a different port | 65 |
| Stop the Web Console | 65 |
| Stop the Web Console from the Start menu | 65 |
| Stop the Web Console from the command line | 65 |
| Install the repository | 66 |
| Upgrade an existing repository | 69 |
| Add Web Console nodes | 70 |
| Use log files | 71 |
| View log files | 71 |
| Edit the obt.log logging properties | 72 |
| Chapter 8: Secure the Web Console | 73 |
| Import a certificate | 73 |
| Implement LDAP | 73 |
| Update webConsole.properties | 74 |
| Include the LDAP manager password | 74 |
| Add users | 75 |

| | |
|---|-----|
| Set password security | 75 |
| Configure HTTPS | 76 |
| Configure HPE IDOL security | 76 |
| Restrict Web services for cloud deployments | 77 |
| Chapter 9: Manage environments | 78 |
| Choose environment type | 78 |
| Standard environments | 78 |
| When to use | 78 |
| Non-intrusive environments | 79 |
| When to use | 79 |
| Use JDBC URLs for database connections | 80 |
| JDBC URLs | 80 |
| Oracle RAC | 81 |
| Configure standard database to file environments for DB2 | 81 |
| Configure parameters | 82 |
| Choose a data movement method for database to database | 82 |
| Configure user-defined business flow parameters | 83 |
| Configure cartridge parameters | 83 |
| Manage locations | 83 |
| Default locations | 84 |
| File archive settings | 84 |
| Change locations and update passwords | 85 |
| Change the archive path | 86 |
| Add locations | 86 |
| Parameterize locations | 87 |
| Configure Hadoop with HPE Structured Data Manager | 88 |
| Refresh test environments | 88 |
| Prerequisites | 88 |
| Rename environments | 94 |
| Chapter 10: Manage business flows | 95 |
| Deploy business flows | 95 |
| Prepare to run business flows | 96 |
| Prepare to reload archived data | 96 |
| Prepare to upload archived data | 97 |
| Recover or cancel business flows and jobs | 97 |
| Access archived data | 98 |
| Create database to database archive access transparency | 98 |
| Relocate database to file archive data | 99 |
| Access the archive data file | 100 |
| Chapter 11: Run business flows and jobs from the command line | 101 |
| Before you begin | 101 |
| Run business flows | 101 |
| Recover and cancel business flows and jobs | 103 |
| Recover business flows | 103 |
| Recover jobs using the recovery script | 103 |

| | |
|--|-----|
| Cancel jobs from the command line | 104 |
| Access the archived data | 105 |
| Create database to database archive transparency | 105 |
| Relocate database to file archive data | 106 |
| Chapter 12: Discovery | 109 |
| HPE SDM Discovery feature | 109 |
| Create a discovery project | 109 |
| Exporting and Importing Projects | 111 |
| Export Projects | 112 |
| Import Projects | 112 |
| Remote Projects | 112 |
| Chapter 13: Map schemas for upload or AQS cache | 113 |
| Map schemas by cartridge | 113 |
| Cartridge mapping files | 115 |
| Map schemas by database | 116 |
| Non-native to native database | 116 |
| Native to non-native database | 116 |
| Non-native to non-native database | 117 |
| Database mapping files | 117 |
| Create and edit database mapping files | 119 |
| Validate the database mapping file | 119 |
| Create mapping by connecting to the source database | 120 |
| Extracting binary zero in Oracle CHAR or VARCHAR | 120 |
| Chapter 14: Perform history schema reconciliation | 122 |
| About the history schema reconciliation process | 122 |
| Overview of the history schema reconciliation process | 124 |
| Prepare to run the history schema reconciliation | 124 |
| Verify current installation | 124 |
| Verify current software version | 125 |
| Verify current patch number | 125 |
| Verify completed jobs and the environment name | 125 |
| Verify necessary changes | 126 |
| Discover differences between the source and target databases | 127 |
| Generate the XML script | 128 |
| Edit the XML script | 129 |
| Deploy the XML script | 132 |
| Verify results | 133 |
| Check the log files | 133 |
| Check the .out file | 133 |
| Check the obt.log file | 134 |
| Edit the logging properties | 134 |
| Chapter 15: Use AQS to access database to file output | 136 |
| Configure the AQS cache | 136 |

| | |
|---|-----|
| Part IV: Appendices | 138 |
| Appendix A: Configuration and runtime parameters | 139 |
| Database to file configuration parameters | 139 |
| Core parameters | 139 |
| Performance parameters | 143 |
| Validation parameters | 146 |
| Centera parameters | 147 |
| Amazon S3 | 148 |
| Upload options | 149 |
| AQS cache options | 149 |
| HPE CM options | 149 |
| Database to database configuration parameters | 150 |
| Core parameters | 151 |
| Performance parameters | 152 |
| Archive access parameters | 156 |
| Advanced selection parameters | 158 |
| Distributed instance parameters | 159 |
| Partition based delete | 161 |
| File naming parameters | 161 |
| File naming parameter usage | 162 |
| Appendix B: Database archiving advanced tasks | 164 |
| SQL tuning | 164 |
| About SQL tuning | 164 |
| About statements | 165 |
| About context values | 165 |
| About statement_name values | 166 |
| About appspack_name and table_identifier values | 168 |
| About using session variable statements | 169 |
| About using custom pre-statement execution code | 169 |
| Edit the sql_tuning.properties file | 170 |
| Specify SQL hints for non-intrusive environments | 174 |
| Specify SQL hints for advanced selection | 175 |
| Gather eligibility analytics for advanced selection | 176 |
| Reclaim space in the active database | 177 |
| Use views for customized reporting | 178 |
| Tune the archive access layer | 178 |
| Implement partition-based deletion for Oracle | 179 |
| Enable SQL trace for Oracle | 180 |
| Enable SQL trace using the sql_tuning.properties file | 180 |
| Add minimal privilege users for Oracle | 181 |
| Lock and unlock the repository | 182 |
| Lock the repository | 183 |
| Unlock the repository | 184 |
| Appendix C: Change encryption keys and passwords | 185 |
| Change the encryption key | 185 |

| | |
|---|-----|
| Reset the encryption key | 185 |
| Change the repository password | 186 |
| Update database location passwords | 187 |
| Appendix D: Scripted product setup and business flow deployment | 189 |
| Understanding properties files | 189 |
| Use special characters in properties files | 189 |
| Script the repository setup | 190 |
| Scripted repository installation procedure | 190 |
| Scripted repository installation properties | 191 |
| Scripted repository installation properties file example | 192 |
| Script the environment creation | 193 |
| Scripted environment creation procedure | 193 |
| Scripted environment creation properties | 193 |
| Scripted environment creation properties file example | 196 |
| Script business flow deployment | 200 |
| Scripted business flow deployment procedure | 200 |
| Scripted business flow deployment properties | 201 |
| Scripted business flow deployment properties file example | 203 |
| Glossary | 205 |
| Index | 211 |
| Send documentation feedback | 216 |

About this document

HPE Structured Data Manager provides powerful tools to build an archive solution that copies or moves data out of your production database and into less expensive storage.

This guide provides information about:

- configuring the Web Console
- deploying business flows
- running business flows and jobs
- monitoring business flows and jobs
- accessing archived data

Intended audience

This guide is intended for:

- users configuring the Web Console
- users running business flows and jobs

Prerequisites

Prerequisites for installing this product include:

- Knowledge of the operating system
- Database knowledge
- Application knowledge

New and revised information

This document includes the following new and revised features in HPE Structured Data Manager:

- Archive Query Server Cache
- Partition-based delete for Oracle
- Enhanced Web Console user interface

See the *HPE Structured Data Manager Installation Guide* for more details about new and revised features for this release.

Related documentation

- *HPE Structured Data Manager API Reference Guide*
Provides information about the Groovy script API files for HPE Structured Data Manager.
- *HPE Structured Data Manager Concepts Guide*
Explains the major concepts of database archiving in general and HPE Structured Data Manager in particular.
- *HPE Structured Data Manager Installation Guide*
Explains how to use the Installer to install the product.
- *HPE Structured Data Manager Tutorial*
Provides step-by-step instructions to build a sample archiving module, deploy it, run it, and troubleshoot errors.
- *HPE Structured Data Manager Developers Guide*
Explains how to use the Designer component to design, build, test, and deploy your archiving projects.
- *HPE Structured Data Manager Runtime Guide*
Explains how to use the Web Console component to run, monitor, and administer business flows that move data to and from the database.
- *HPE Structured Data Manager Troubleshooting Guide*
Explains how to diagnose and resolve errors, and provides a list of common errors and solutions.
- *HPE Structured Data Manager Upgrade Guide*
Explains how to upgrade the product and upgrade the archive schema generated by earlier versions of the product.
- *HPE Structured Data Manager Release Notes*
Lists any items of importance that were not captured in the regular documentation.
- *HPE Structured Data Manager PeopleSoft Modules Installation and Deployment Guide*
Explains how to install the PeopleSoft integration kit.
- *HPE Structured Data Manager Oracle E-Business Suite Modules Installation and Deployment Guide*
Explains how to install the Oracle E-Business Suite integration kit.

Part I: Introduction

This part provides an introduction to the basics of database archiving. It includes:

- [Introduction to structured data management](#)

Chapter 1: Introduction to structured data management

Managing structured data is the act of moving or retiring data from an active database to an archive data store. An active database is typically an online transaction processing (OLTP) database and is also sometimes referred to as a production database. An archive data store is another database (sometimes called an archive database) or a set of files (CSV or XML). *Moving* means copying the specified parts of the active database to the archive data store and deleting it from the active database.

Optimization is desirable for many reasons, but two of the most common ones are:

- **To retain data for the long term while keeping it accessible for the purposes of corporate governance, eDiscovery, and/or legal hold.** Database to file movement enables you to move or copy large amounts of data from a database to structured files.
- XML files provide an open, standards-based format that can be opened long after the original application has been retired. As such, it provides an excellent method of storing data for long term retention.
- Comma separated values (CSV) and JavaScript Object Notation (JSON) files provide a simple format that many applications can import.
- **To reduce the footprint of your active database.** After databases and their associated applications have been in use for some time, they often become very large, which slows performance and increases the need for additional, expensive hardware. Archiving older, infrequently accessed data from the active database enables you to move eligible data to less expensive storage while also reorganizing the active database. This reduction in footprint tends to improve the performance of the active database.
- [Before you begin](#)
- [Web Console overview](#)

For a complete conceptual introduction to HPE Structured Data Manager, see the *HPE Structured Data Manager Concepts Guide*.

Before you begin

Before you begin performing the tasks in this guide, you should:

1. Review the *HPE Structured Data Manager Concepts Guide* to become familiar with the software and how you plan to use it.
2. Install the HPE Structured Data Manager software according to the instructions in the *HPE Structured Data Manager Installation Guide*.
3. Go through *HPE Structured Data Manager Tutorial*. The tutorial enables you to get hands on with the product quickly and exposes you to many of the most commonly used features.
4. Use the *HPE Structured Data Manager Developer's Guide* to design and develop or customize your archive solution.

Web Console overview

Most of the work of archiving data is performed in the Web Console. The Web Console is a powerful browser-based interface used to:

- Configure the repository
- Create environments
- Create and modify users
- Deploy business flows
- Launch business flows to move data
- Monitor running business flows and their history

The Web Console enables you to remotely manage multiple environments for the movement of data.

Part II: Tutorials

This part provides a set of tutorials that build on the basic tutorial described in *HPE Structured Data Manager Tutorial*. It includes:

- [Configure the Web Console](#)
- [Deploy and run business flows](#)
- [Upload to HPE Vertica](#)
- [Perform multiple data movements](#)
- [View eligibility analytics](#)

Chapter 2: Configure the Web Console

Before you can deploy, run, and monitor your business flows to move data, you must configure the Web Console. The Web Console is the primary runtime interface to HPE Structured Data Manager.

This tutorial walks you through the process of configuring the Web Console, including:

NOTE: Before you begin this tutorial, make sure that you have installed HPE Structured Data Manager according to the *HPE Structured Data Manager Installation Guide*.

- [Start the Web Console](#)
- [Install the repository](#)
- [Create an environment](#)
- [Configure settings](#)
- [Manage the home directory](#)
- [Create and edit Web Console users and groups \(optional\)](#)

Start the Web Console

The Web Console must be running before you can connect to it in your browser. After you have started the Web Console, you can perform operations such as installing the repository, creating environments, and managing users.

NOTE: When you install HPE Structured Data Manager, the Installer gives you options to start the Web Console and launch its interface in your browser when the installation completes. If you chose those options and the Web Console is already running, you need not perform the steps in this section.

TIP:

If the default ports used by the Web Console are already in use by another application, you need to start the Web Console with other ports. The process for changing ports is done through Tomcat. See [Start and configure the Web Console, on page 63](#).

To start the Web Console

1. Ensure that you are logged into your operating system as a user with appropriate administrative privileges.
2. Start the Web Console in one of the following ways:

From program group (Windows)

- a. On Microsoft Windows, select **Start > HPE Structured Data Manager > Start Web Console**.

Or:

From command line

- a. Open a command window. For example, on Microsoft Windows, select **Start > Run**, enter `cmd`, and click **OK**.
- b. Change to the `bin` directory where you installed HPE Structured Data Manager. For example, on Microsoft Windows:

```
cd c:\Program Files\HPSDM730\obt\bin
```

On UNIX:

```
cd /home/HPSDM730/obt/bin
```

- c. Enter the start command:

On Windows:

```
webConsole.bat start
```

On UNIX:

```
./webConsole.sh start
```

HPE Structured Data Manager uses Tomcat to run the Web Console. On Windows, a new window is spawned for the Tomcat process. On Linux, a Tomcat process is started.

TIP: The first time you launch the Web Console, HPE Structured Data Manager must perform some initial configuration, which may take longer. You should not interrupt this process. If the Web Console fails to launch or you inadvertently interrupt the initiation process, see the *HPE Structured Data Manager Troubleshooting Guide*.

Check the Tomcat process and logs to determine if the Web Console started successfully.

From Services (Windows)

1. On Microsoft Windows, install the Web Console service, if you did not install it from the Installer:

```
<install_dir>\obt\bin\WebConsoleService.bat install
```

2. Start the service:

```
<install_dir>\obt\bin\WebConsoleService.bat
```

TIP:

You can also use `WebConsoleService.bat` to stop and uninstall the service.

NOTE: You can also start and stop the HPE Structured Data Manager service from **Control Panel > Administrative Tools > Services**. By default, the service is configured to start automatically during system start. You can also start it manually. Ensure that the user specified to run the service has appropriate administrative privileges on the system.

3. Launch a supported browser, and connect to the Web Console using the following URL:

```
http://localhost:8080/WebConsole
```

The first time that you start the Web Console after installing HPE Structured Data Manager, you must install a new repository or import one from an earlier version. For information on installing a new repository, see [Install the repository, on the next page](#). For information on upgrading an existing repository, see the *HPE Structured Data Manager Upgrade Guide*.

Install the repository


HPE Structured Data Manager stores metadata in a repository to facilitate the deployment and running of business flows, and administering the Web Console. You can create the repository in an Oracle, SQL Server, or PostgreSQL database.

NOTE: For information about database prerequisites for the repository, see the *HPE Structured Data Manager Installation Guide* and the HPE Structured Data Manager support matrix.

To install the repository

1. Launch the Web Console in your browser.

The Repository installation screen displays:



2. Click **Install a New Repository**. The Repository Database: Administrator page displays.
3. Select the repository type from the **Type** list:
 - **Oracle** means use an Oracle database and connect to it using the basic Oracle connection method.
 - **SQL Server** means use a SQL Server database.
 - **PostgreSQL** means use a PostgreSQL database.

NOTE: You can only have one HPE Structured Data Manager repository per PostgreSQL database. If necessary, create another database, which can be on the same server.

- **JDBC URL** means connect to a database using a JDBC URL string. For example, to connect to an Oracle database with Real Application Clusters (RAC) you must use JDBC URL.

When you select a repository type, the properties of the Administrator page change accordingly. For example, when you choose SQL Server, a DB Server field is added, the default user changes to sa, and the default port changes to 1433.

For Oracle

An example for Oracle:

Repository Database: Administrator

The administrator account is used by the deployment assistant to create users and grant access to resources.

| | |
|------------------|---------------|
| Repository Type: | Oracle ▼ |
| User:* | system |
| Password:* | ***** |
| Host:* | 16.103.11.102 |
| Port:* | 1521 |
| Service Name:* | worcl1 |

For SQL Server

An example for SQL Server:

Repository Database: Administrator

The administrator account is used by the deployment assistant to create users and grant access to resources.

| | |
|------------------|---------------|
| Repository Type: | SQL Server ▼ |
| User:* | sa |
| Password:* | ***** |
| Host:* | 16.103.11.213 |
| Port:* | 5002 |
| Instance Name:* | SQLSERV12 |

For PostgreSQL

An example for PostgreSQL:

Repository Database: Administrator

The administrator account is used by the deployment assistant to create users and grant access to resources.

| | |
|------------------|--------------|
| Repository Type: | PostgreSQL ▼ |
| User:* | postgres |
| Password:* | ***** |
| Host:* | localhost |
| Port:* | 5432 |
| Database Name:* | postgres |

4. Enter the credentials appropriate for the selected database type.

| Field | Description |
|---------------|--|
| User | Enter the database administrator username. |
| Password | Enter the database administrator password. |
| Host | (Oracle, SQL Server, PostgreSQL) Enter the name or address of the machine where the database is installed. |
| Port | (Oracle, SQL Server, PostgreSQL) Enter the port number of your Oracle, SQL Server, or PostgreSQL database. |
| Service Name | (Oracle only) Enter the name of your Oracle database, for example, ORCL. |
| DB Server | (SQL Server only) Enter the name of the SQL Server database, for example, MSOLTP. |
| Database Name | (PostgreSQL only) Enter the name of the PostgreSQL database, for example, sdm_repos. You can only have one repository per PostgreSQL database. |

5. Click **Next**.

The Repository Database: User page displays.

For Oracle

Repository Database: User

Create the repository user. Repository user is a database user for logging in to the repository schema.

| | |
|--------------------|--|
| User:* | <input type="text" value="obt_rep"/> |
| Password:* | <input type="password" value="Password"/> |
| Confirm Password:* | <input type="password" value="Confirm Password"/> |
| Data Tablespace: | <input type="text" value="USERS (45 MB)"/> <input type="button" value="v"/> |
| Temp Tablespace: | <input type="text" value="TEMP (< 1 MB)"/> <input type="button" value="v"/> |

Repository Database: Encryption Key

The Encryption Key is used as Repository Master Password

| | |
|--------------------------|---|
| Encryption Key:* | <input type="password" value="Encryption Key"/> |
| Confirm Encryption Key:* | <input type="password" value="Confirm Encryption Key"/> |

Directories

| | |
|--------------------------|---|
| Home Directory:* | <input type="text" value="C:\ProgramData\Hewlett-Packard\HPOBT"/> |
| Base Archive Directory:* | <input type="text" value="C:\ProgramData\Hewlett-Packard\HPOBT\archive"/> |

* Required Fields

« Back **Next »**

For SQL Server

Repository Database: User

Create the repository user. Repository user is a database user for logging in to the repository schema.

User:*

obt_rep

Password:*

Password

Confirm Password:*

Confirm Password

DB Name:*

obt_rep

Primary data Size:*

90

Transaction Log Size:*

50

Repository Database: Encryption Key

The Encryption Key is used as Repository Master Password

Encryption Key:*

Encryption Key

Confirm Encryption Key:*

Confirm Encryption Key

Directories

Home Directory:*

C:\ProgramData\Hewlett-Packard\HPOBT

Base Archive Directory:*

C:\ProgramData\Hewlett-Packard\HPOBT\archive

* Required Fields

« Back

Next »

For PostgreSQL

Repository Database: User

Create the repository user. Repository user is a database user for logging in to the repository schema.

User:*

obtrep

Password:*

Password

Confirm Password:*

Confirm Password

Repository Database: Encryption Key

The Encryption Key is used as Repository Master Password

Encryption Key:*

Encryption Key

Confirm Encryption Key:*

Confirm Encryption Key

Directories

Home Directory:*

C:\ProgramData\Hewlett-Packard\HPOBT

Base Archive Directory:*

C:\ProgramData\Hewlett-Packard\HPOBT\archive

* Required Fields

« Back

Next »

6. Select or enter the following information for the repository user:

| Field | Description |
|------------------------|---|
| User | Enter the name you want to use for the repository user. The default value is obt_rep. |
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Data Tablespace | (Oracle only) Select the data tablespace you want to use. |
| Temp Tablespace | (Oracle only) Select the temporary tablespace you want to use. |
| DB Name | (SQL Server only) Enter the SQL Server database name. |
| Primary Data Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the data file. |
| Transaction Log Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the log file. |
| Encryption Key | Enter the encryption key you want to use. |
| Confirm Encryption Key | Enter the encryption key again. |
| Home Directory | Modify the home directory for HPE Structured Data Manager or accept the default value. |
| Base Archive Directory | Modify the directory where archive data is stored or accept the default value. |

7. Click **Next**.

The Console Administrator page displays.

Console Administrator

The Console Administrator is initially the only user authorized to connect to the Web Console.

| | |
|--------------------|-----------------------|
| User Name:* | admin |
| Password:* | •••••• |
| Confirm Password:* | Confirm Password |
| Real Name: | Users display name |
| Description: | User Description |
| Email: | Contact email address |
| Phone Number: | Contact phone number |

* Required Fields

« Back Next »

8. Select or enter the following information:

NOTE: By default, this user is assigned to the Administrators group, which has all privileges enabled.

| Field | Description |
|------------------|---|
| User Name | Enter the name you want to use for the Web Console administrator. The Web Console administrator is used to log on to the Web Console. |
| Password | Enter the password for the administrator. |
| Confirm Password | Enter the password again. |
| Real Name | Optionally, enter the name of the administrator user. |
| Description | Optionally, enter a description. |
| Email | Optionally, enter an email address. |
| Phone Number | Optionally, enter a phone number. |

9. Click **Next**. The Summary page appears for your review.
10. Click **Finish**.

It may take several moments for the repository creation to complete. When the repository is created, the following message displays:

You have successfully installed the repository. Please do NOT close your browser. Web Console will restart.

After the Web Console restarts, the login screen displays.

TIP: If for some reason the repository installation should fail, you need to stop and start the Web Console before trying again.

11. Enter the Web Console administrator user name and password that you just created.
12. Click **Login**.

You are now logged in to the Web Console and can create a new environment.

Create an environment

Before you can deploy your business flows, you must create a deployment environment for them with the appropriate characteristics, such as the source database credentials and available data movement methods.

To create an environment

1. Launch the Web Console in your browser and login as the administrative user you created when you installed the repository, [Install the repository, on page 19](#).
2. Click **Business Flow Management** from the Web Console home page.
3. If you have not previously created an environment, the Web Console automatically prompts you to

create one when you open it. Otherwise, click **Environment** at the top of the page and then **New** in the left navigation pane to display the New Environment wizard.

4. Specify a name for the new environment, for example, `Oracle_OLTP` or `SQLServer_OLTP`.
5. Specify whether you want a non-intrusive environment configuration. For the purposes of this tutorial, leave **Non-intrusive** unselected. For more information, see [Manage environments, on page 78](#).
6. Choose a user group from the **Authorized Group** list to which you want to grant access to the environment. Members of this group will have privileges on the environment.

TIP: If you want to give other groups privileges on the environment, change the group membership, or modify the privileges of the group, see [Create or edit a user group, on page 40](#).

7. Choose the user group to which you want to grant access to the user group. See [Create and edit Web Console users and groups \(optional\), on page 40](#) for more information.
8. Select the source database type.
 - **Oracle**
 - **SQL Server**
 - **Sybase**
 - **JDBC URL**
 - **DB2**

TIP: For SQL Server, ensure that the SQL Server Browser service is running and you are using a fixed port rather than a dynamic one. See [Manage the home directory, on page 39](#).

Selecting **JDBC URL** indicates that you want to connect to a database using a JDBC URL string. A common use of the JDBC URL is for generic JDBC non-intrusive connections. One required usage is to connect to an Oracle database with Real Application Clusters (RAC).

9. Enter the necessary credentials for the source database. For more information, see the Web Console online help system.

TIP: If you chose **JDBC URL** for database type, you should not choose the **Remember** check box.

For Oracle

New Environment

Environment Name *
Oracle_OLTP

Description

Non-Intrusive
☐

Authorized Group
Administrators

Source Database

Source Database Type *
Oracle

Administrative User:
system

Password *
.....

Remember
☒

Host *

Port *
1521

Service Name:

For SQL Server

New Environment

Environment Name *
SQLServer_OLTP

Description
SQLServer_OLTP

Non-Intrusive
☐

Authorized Group
Administrators

Source Database

Source Database Type *
SQL Server

Administrative User:
sa

Password *
.....

Remember
☒

Host *
16.103.11.213

Port *
5002

Instance Name:

NOTE:
For SQL Server, you can set **Windows Authentication** to **true**, in which case HPE Structured Data Manager uses your operating system user to connect to SQL Server. If you choose Windows Authentication, you only need to fill in the Host, Port, and Instance Name.

For DB2

New Environment

Environment Name * DB2_OLTP_01

Description DB2_OLTP

Non-Intrusive ☐

Authorized Group Administrators ☒

Source Database

Source Database Type * DB2 ☒

Administrative User: db2/97a

Password * *****

Remember ☒

Host * hoye.svs.usa.hp.com

Port * 50000

Database: QA_USER

10. Click **Next**.

The database to database option displays. Optionally, you can add database to database movement. For the purposes of this example, you will configure database to database.

11. Select **Include Database to Database Capability** to include database to database movement in the environment. When you select Include Database to Database Capability, the options for single and distributed instances display.

NOTE: For DB2, the tablespace page size for your history database must accommodate additional columns added to the history tables. The table space page size should be greater than or equal to the table space on the source database for all managed tables.

NOTE: For DB2, complex (structured) UDTs are not supported.

12. If there is a separate database for archived data, click the **Archive Data in another database (Distributed Instance)** radio button. The fields for a target database type and credentials appear. Otherwise, select **Archive Data within the same database (Single Instance)**.

NOTE: For distributed instances on DB2, you must enable federation on both source and target databases and enable the TWO_PHASE_COMMIT property. You also need to create a user mapping between the interface and relocation user so that each can access the federation (dblink). Furthermore, the archive access user and the relocation user must also be mapped for the federation. See your DB2 documentation for more information about setting up federation.

13. If necessary, enter the information for your target database.

For Oracle

Database to Database Capability

☒ Include Database to Database Capability

Configuration

☐ Single Instance (Archive Data within the same database)

☒ Distributed Instance (Archive Data in another database)

Target Database

Administrative User *

system

Password *

••••••

Host *

16.103.11.102

Port *

1521

Service Name *

WPDB2

Source TNS Name

Target TNS Name

Source to History dblink *

OBT_OLTP_TO_HIST

History to Source dblink *

OBT_HIST_TO_OLTP

For SQL Server

Database to Database Capability

☒ Include Database to Database Capability

Configuration

☐ Single Instance (Archive Data within the same database)

☒ Distributed Instance (Archive Data in another database)

Target Database

Administrative User *

sa

Password *

••••••

Host *

16.103.11.213

Port *

5002

Instance Name *

SQLSERV12

For DB2:

« Previous

Next »

Database to Database Capability

☒ Include Database to Database Capability

Configuration

☐ Single Instance (Archive Data within the same database)

☒ Distributed Instance (Archive Data in another database)

Target Database

Administrative User *

db2i97a

Password *

Host *

hoya.svs.usa.hp.com

Port *

50000

Database *

QA_USER

Data Source for History on OLTP *

OBT_OLTP_TO_HIST

Data Source for OLTP on History *

OBT_HIST_TO_OLTP

14. Click **Next**.
The Interface User page displays.

For Oracle

Interface User

The Interface User is created in the Source Database to store intermediate results.

User *

obt_if

Password *

Confirm Password *

Data Tablespace

USERS (20 MB)

▼

Temp Tablespace

TEMP (< 1 MB)

▼

For SQL Server

Interface User

The Interface User is created in the Source Database to store intermediate results.

User *

obt_if

Password *

Confirm Password *

DB Name *

obt_if

Primary data Size*

50

Primary data Size*

25

For DB2:

Interface User

The Interface User is created in the Source Database to store intermediate results.

User *

TESTER

✓

Password *

Confirm Password *

Schema *

obt_if

Data Tablespace

USERSPACE1 (13 MB)

✓

15. Enter the following values for the interface user:

NOTE: For DB2, you should pre-create an interface user to select here. The user cannot be the same as your administrative user. The user must have CONNECT privileges on the source database. See your DB2 documentation for information about creating users and setting user privileges.

| Field | Description |
|----------------------|---|
| User | For Oracle and SQL Server, enter the name you want to use for the interface user. The default value is obt_if. For DB2, choose an existing operating system user from the list of values. The list of values is populated with users that have connect privileges on the DB2 database. |
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Schema | (DB2 only) Enter the name of the schema you want to use for the interface user. The default value is obt_if. |
| Data Tablespace | (Oracle and DB2) Select the data tablespace you want to use. |
| Temp Tablespace | (Oracle only) Select the temporary tablespace you want to use. |
| Primary Data Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the data file. |
| Transaction Log Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the log file. |

16. Click **Next**.
17. If you selected distributed archive in [If necessary, enter the information for your target database., on page 27](#), the Relocation User page displays. Otherwise, you can skip to [Enter the following information for the user who builds archive access:, on page 32](#).

For Oracle

Relocation User

The Relocation User is used to transfer the data.

| | |
|--------------------|---|
| User * | <input type="text" value="obt_reloc"/> |
| Password * | <input type="password" value="....."/> |
| Confirm Password * | <input type="password" value="....."/> |
| Data Tablespace | <input type="text" value="USERS (4 MB)"/> ▼ |
| Temp Tablespace | <input type="text" value="TEMP (< 1 MB)"/> ▼ |

For SQL Server

Relocation User

The Relocation User is used to transfer the data.

User *

obt_reloc

Password *

Confirm Password *

DB Name *

obt_reloc

Primary data Size*

50

Primary data Size*

25

For DB2

Relocation User

The Relocation User is used to transfer the data.

User *

TESTER

Password *

Confirm Password *

Schema *

obt_reloc

Data Tablespace

USERSPACE1 (13 MB)

18. Enter the following values for the relocation user:

NOTE: For distributed instances with DB2, your relocation user must have CONNECT privileges on the target database. See your DB2 documentation for information about setting user privileges.

| Field | Description |
|----------------------|---|
| User | For Oracle and SQL Server, enter the name you want to use for the relocation user. The default value is obt_reloc. For DB2, choose an existing operating system user from the list of values. The list of values is populated with users that have connect privileges on the DB2 database. |
| Password | Type the password for the user you created. |
| Confirm Password | Type the password again. |
| Schema | (DB2 only) Enter the name of the schema you want to use for the relocation user. The default value is obt_reloc. |
| Data Tablespace | (Oracle and DB2) Select the data tablespace you want to use. |
| Temp Tablespace | (Oracle only) Select the temporary tablespace you want to use. |
| Primary Data Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the data file. |
| Transaction Log Size | (SQL Server only) Accept the default value or enter a maximum size in MB for the log file. |

The Archive Access page displays.

For Oracle

Archive Access

Create the archive access user. This account is used only for building archive access

Locations

Source

User *

obt_build_aa

Password *

.....

Confirm Password *

.....

Data Tablespace

USERS (20 MB)

Temp Tablespace

TEMP (< 1 MB)

For SQL Server

Archive Access

Create the archive access user. This account is used only for building archive access

Locations

Source

User *

obt_build_aa

Password *

.....

Confirm Password *

.....

For DB2

Archive Access

Create the archive access user. This account is used only for building archive access

Locations

Source

User *

TESTER

Password *

.....

Schema *

obt_build_aa

Data Tablespace

USERSPACE1 (13 MB)

Data Tablespace

USERSPACE1 (13 MB)

History Data Source for AA on OLTP *

OBT_OLTP_TO_HIST

19. Enter the following information for the user who builds archive access:

NOTE: For DB2, your archive access user must have CONNECT privileges on the source database. See your DB2 documentation for information about setting user privileges.

| Field | Description |
|------------------|--|
| Location | Select Source or Target. |
| User | For Oracle and SQL Server, enter the name you want to use for the archive access user. The default value is obt_build_aa. For DB2, choose an existing operating system user from the list of values. The list of values is populated with users that have connect privileges on the DB2 database. |
| Password | Type the password for the user you created. |
| Confirm Password | Type the password again. |

| Field | Description |
|------------------------------------|---|
| Schema | (DB2 only) Enter the name of the schema you want to use for the archive access user. The default value is obt_build_aa. |
| Data Tablespace | (Oracle only) Select the data tablespace you want to use. |
| History Data Source for AA on OLTP | (DB2 only) Enter the name of the history data source used for archive access. The default value is OBT_OLTP_TO_HIST. |
| Temp Tablespace | (Oracle only) Select the temporary tablespace you want to use. |

20. Click **Next**.

The summary page for the environment opens. Review the settings that you have chosen.

21. Click **Finish** to create the environment.

It may take several moments to create the environment. When the environment is ready, the success message should appear at the top of the page.

For the purposes of this tutorial, you can now optionally create a Web Console user [Create and edit Web Console users and groups \(optional\)](#), on page 40, or skip directly to [Deploy and run business flows](#), on page 43.

Configure settings

The Web Console includes a number of other settings that you can configure from the **Settings** menu.

For more information on these settings, see the Web Console Help system.

Set up email access

You can set up the Web Console to send out email alerts for various events. Before you can send email using the Web Console, you need to configure your Server Email settings.

To configure the settings

1. Navigate to **Settings** in the Web Console.
2. Click **Email Server**.
The Email Settings page opens.
3. To require authentication (password), select **Authentication Required**.
4. Enter the email settings:
 - Web Server URL: The URL for the server where the Web Console is running. This is the URL that other machines need to connect to the Web Console, for example, `http://<host>:8080/WebConsole`.
 - Email Host: The host name of the mail server you want the Web Console to use for sending messages, for example, `smtp.gmail.com`.
 - Email Port: The mail server port number, for example, 465.

- Email User Name: The email account name, for example, tom.baker@hiscompany.com.
 - Email Password: The password to the email account. This field is visible only if you selected **Authentication Required**.
 - Email From: The email address that will appear in the From part of the email. It might be the same as User Name, but you could also enter a different name here.
 - Additional Settings: Enter any additional settings needed to configure email, such as SMTP socketFactory settings. For example:
mail.smtps.starttls.enable=true
mail.smtp.socketFactory.fallback=false
mail.smtp.socketFactory.class= javax.net.ssl.SSLSocketFactory
3. Click **Apply**.
 4. To test whether your server settings are correct and working, navigate to the following URL:
http://<host>:8080/WebConsole/main/mailTest

NOTE:

You must be an administrator for this page to work for you, and your email information must be set up at <http://<host>:8080/WebConsole/settings/info>.

Manage indexing server settings

On the Indexing Server Settings page, you can specify the indexing server and options to use when running indexing cartridges.

To manage indexing server settings

1. Navigate to **Settings** in the Web Console.
2. Click **Indexing Server**.
3. Choose the **Indexing Server Type**, IDOL or Apache Solr. If you choose Apache Solr, you will only see a subset of the settings that follow.
4. Select the protocol to employ when communicating with the IDOL server, **HTTP** or **HTTPS**. See the *HPE IDOL Administration Guide*, "Set Up an SSL Connection," for more information about HTTPS and IDOL.
5. Enter the **Server Host** where IDOL is located.
6. Enter the **Server Action Port** and **Server Content Port** for IDOL, for example 9000 and 9001.
7. If you wish to enable TEA encryption, check **Enabled** and enter your TEA encryption key in the provided field. For more information about TEA, see the *HPE Structured Data Manager Developer's Guide*.
8. Enter the **Default IDOL Database Name**. You can enter a different database name in the cartridge editor for each cartridge to override this default.
9. Choose a database security option:

- **Allow users with data access privileges to query any IDOL database** means that any user who has the data access privilege granted to them can also query any IDOL database.
 - **Grant users IDOL database query access** means that you must explicitly grant users query access to IDOL databases. If you choose this option, when you select **Apply**, a new link appears called **Set security by database**. Following the link takes you to a page where you can manage query access by IDOL database and user.
10. Optionally, modify the Groovy script that specifies the entry ID.
 11. Click **Apply**.

Configure the trusted store for certificates

For locations that require secured communications, you need to store the certificates from trusted Certificate Authorities (CAs).

To store the certificates

1. Navigate to **Settings** in the Web Console.
2. Click **Certificate Issuers**.
3. Enter the certificate settings.
 - **File Path:** The file that contains all of the certificates from CAs that are trusted. Any certificate not issued by one of these issuers will not be accepted by the HTTPS processing. Enter the full path name to this file.
 - **Password:** The password for the trust store.
 - **Type:** The type of trust store (JKS, PKCS12, and so on).
4. Click **Apply**.

Masking Server settings

HPE Structured Data Manager now supports the Format Preserving Encryption (FPE) and Secure Stateless Tokenization (SST) provided by HPE SecureData Enterprise. The Masking Server Settings page includes a default public instance so you can try FPE/SST. However, HPE recommends that you purchase the on-premises HPE SecureData Enterprise product.

Prerequisites

The HPE SecureData SimpleAPI (part of HPE SecureData Enterprise) version 5.1 or later must be installed on your server.

On the **Masking Server Settings** page, you can specify the masking server and options.

To manage masking server settings

1. Navigate to **Settings** in the Web Console.
2. Click **Masking Server**.
The HPE SecureData Appliance Settings page opens.
3. Specify the HPE SecureData Appliance settings:
 - a. **Masking Product**.
 - i. **HPE SecureData Enterprise - FPE**. Indicates that HPE SDM will use HPE SecureData Enterprise Format Preserving Encryption (FPE) as the masking server.
 - b. **Shared Library/DLL Path**. Specify the path to `vibesimplejava.dll`, based on your installation of HPE SecureData API installation directory.
 - c. **Server Base URL**. Specify the base URL of the HPE SecureData Appliance.
protocol://<serverDNS:Port>
where
protocol is HTTP or HTTPS
serverDNS is the server DNS or IP address
Port is the port number of the HPE SecureData Appliance.

Example

`http://myHPESecureDataServer.myDomain.com` or

`https://myHPESecureDataServer.myDomain.com`

NOTE:

The default value of **Server Base URL** points to the public instance of HPE SecureData.

- d. **Default Identity**. Specify the default identity value used, if the per format identity is not provided. This specified identity value is also used for AUTO format while testing the connection.
- e. Enter the data to be masked or unmasked in the **Input Data** box.
This is an optional field, not required if you use the **Test Connection**, but required test the masking or unmasking.
- f. Click **Test Connection** to test the connection to the masking server.
- g. Click **Test Mask** to test the data masking.
- h. Click **Test Unmask** to test the data unmasking.
The Output Data box displays the masking or unmasking results.

- i. Optionally, click **Clear** to clear the values in the **Input Data** box.
 - ii. Click the arrow button to copy the value from the **Output Data** box to the Input Data box.
This allows you to quickly verify the masking and unmasking feature.
4. Click **Apply**.

NOTE:

The **Apply** button is enabled only if the connection test succeeds.

The HPE SecureData Format - SDM Functions page opens. This page allows you to enter the display names for the FPE/SDM functions and also for the per format identity.

5. Enter a function prefix in the **Function Prefix** box.

NOTE:

The special characters <, >, :, ", /, \, |, ?, and * and blanks are not allowed in the **Function Prefix** box.

6. Click **Generate** to generate the functions.

Once you generate the functions, the page status bar displays shows the number of functions generated and the directory name where the Groovy files are stored.

Use this location to import the FPE/SST functions into HPE SDM Designer.

Date-Time Formats in SDM

In HPE Structured Data Manager database, the data types for Date, DateTime or Timestamp are treated in the same format as Timestamp type. The internal string representation of the Timestamp is of the format YYYY-MM-DD HH24:MI:SS.nnnnnnnnn (where n representation nanoseconds part). However, when you use the same with SecureData for encryption or decryption the nanosecond part is ignored.

SecureData appliance has ORA-DATE format which is different from the SDM internal representation. You must create another Date type format in SDA to ensure that the Date-Time format works correctly. You can create the same with the format string as YYYY-MM-DD HH24:MI:SS.

View Date format

| Format Name and Value | |
|-----------------------|--------------------------------------|
| Format Name | SDM-ORA-DATE |
| Format String | YYYY-MM-DD HH24:MI:SS |
| Minimum Year | 1900 |
| Maximum Year | 2017 |
| Data Protection Type | FPE - Format-Preserving Encryption ▼ |
| Comment | Date between years 1900 and 2017 |

NOTE:

Ensure to restart the SDM Web Console whenever you add a new format to SDA.

Data Access Cartridge

If the data is to be identified or decrypted, then the de-identified data column requires the pseudo column. The identified data is available to the user based on the privileges set in SDA.

In the pseudo column one of the following groovy functions is used to identify the data:

```
String access(String format, String dataIn)

Timestamp access(String format, Timestamp dataIn)

BigDecimal access(String format, BigDecimal dataIn)
```

The format string is to be specified based on the format of the data stored in the column and appropriate conversion function is to be used for dataIn.


For each of the function here are the example usages:

```
secureData.access('AUTO', CUSTOMER.EMAIL.toString())

secureData.access('AUTO', CUSTOMER.DATEOFBIRTH.toTimestamp())

secureData.access('AUTO', CUSTOMER.CUSTOMERID.toBigDecimal())
```

For example:



Similar identification of the encrypted data can be used to protect the method available in secureData object and to encrypt the data in DAC. You can use the following methods to do this:

```
String protect(String format, String dataIn)

Timestamp protect(String format, Timestamp dataIn)

BigDecimal protect(String format, BigDecimal dataIn)
```

Manage the home directory

HPE Structured Data Manager stores its program data, such as configuration, log, and archive files, in a location that is separate from its program files. This home directory is determined at installation time:

- By default, HPE Structured Data Manager attempts to configure its home directory in the standard location for application data on the operating system. For example, on Windows, the Installer tries to use C:\ProgramData.
- By default, HPE Structured Data Manager attempts to configure its home directory in the standard location for application data on the operating system. For example, on Windows, the Installer tries to use C:\ProgramData.

After installation, you can change the home directory by following the instructions in [Change the home directory, below](#).

Example

Configuration files might be stored in:

```
<app_data_dir>/Hewlett-Packard/HPOBT/config
```

where <app_data_dir> is the location where your operating system stores application data, or your own custom location.

TIP: On Windows, the ProgramData directory may be hidden by default. You need to show that directory in order to find the log files. See your MS Windows documentation for information on displaying hidden folders and files.

Change the home directory

To change the home directory for HPE Structured Data Manager:

1. Stop all HPE Structured Data Manager processes, including Web Console and AQS.
2. Open <install_dir>\obt\config\obt.env in a text editor and change existing path to the new path.

For example:

```
OBT_HOME=C:/SDM_Home/HPOBT
```

NOTE: Notice that, even on Windows, the directory separator is a forward slash (/).

3. In your file system, copy all of the directories and files except aqsdatasources from your current home directory to the new one that you just specified in obt.env.
4. For aqsdatasources, run the following command from <install_dir>\obt\bin:
oacommandexec -m <data src name> <new location>

For example:

```
oacommandexec -m xmlArchive D:\home1\aqsdatasources\xmlArchive
```

5. Restart HPE Structured Data Manager processes, such as Web Console and AQS.

Create and edit Web Console users and groups (optional)

In most cases, more than one user needs to perform functions in the Web Console, but not every user needs the same privileges. For example, some users might only need to run jobs while others may need to deploy business flows, set parameter values, and run jobs. Web Console provides a user model that enables you to create as many users as you need, and grant or withhold privileges. For cases where you have a many users to manage, you can also create groups to make assigning privileges and access easier.

- [Create or edit a user group, below](#)
- [Create or edit a user, on the next page](#)

Create or edit a user group

Groups are used to assign privileges and environments to users.

NOTE: Typically, you should have at least one group that has all privileges enabled and has at least one member. By default for new installations, HPE Structured Data Manager creates such a group for you and names it Administrators. The user you created when you installed the repository is assigned to this group. Similarly, for new default installations, HPE Structured Data Manager creates a group called Users that has the necessary privileges to run business flows and change parameters.

To create or edit a user group

1. Launch the Web Console in your browser and log in.
2. Click **Business Flow Management** from the Web Console home page.
3. Click **Users > Groups**. The Groups page opens.

If you have not previously created any other groups, you should only see the Administrators and Users groups created by default during initial configuration Web Console.

4. Click **New Group** or click on an existing group.

The Create/Edit Group page opens.

5. Enter or modify a **Group Name**.

NOTE: For an LDAP group, use underscores in place of spaces in the name.

6. Use the **Add** and **Remove** buttons to specify which of the available privileges you want to assign to this group. For the purposes of this example, you do not need to make any changes.
7. Click **Members**.
8. Use the **Add** and **Remove** buttons to select which users to make members of this group. For the

purposes of this example, make the user you just created in [Create or edit a user](#), below a member of this group.

9. Click **Environments**.
10. Use the **Add** and **Remove** buttons to select the environments to which this group has access. For the purposes of this example, give the group access to the environment you created in [Create an environment, on page 24](#).
11. Click **Save**.

Create or edit a user

To create or edit a user

1. Launch the Web Console in your browser and login.
2. Click **Business Flow Management** from the Web Console home page.
3. Click **Users > Manage Users**.
The User List page displays. If you have not previously created any other users, you should only see the user you created when you installed the repository, for example, `admin`.
4. Click **New** or an existing user.
The Create/Edit User page displays.
5. If LDAP is enabled, you see the **Type** field on this page. If you see it, choose **Built-in** for the purposes of this tutorial. Otherwise, you can skip this step.

TIP: Even if you have LDAP configured, you may want to maintain at least one built-in user that you can use to login should your LDAP directory go down for any reason.

6. Enter at least the following values:
 - **Login Name:** tom_baker
 - **Full Name:** Tom Baker
 - **Password:** welcome
 - **Enabled:** checked
7. Use the **Add** and **Remove** buttons to specify to which of the available groups this user should belong. For the purposes of this example, you do not need to make any changes.
8. Click **Create** at the top of the page.
The Show User page opens.

Verify the new Web Console user

To verify the Web Console user

1. Click **Logout**.
2. Log back in as the user you just created.
3. Click **Business Flow Management**.
Notice how the User link does not appear to Tom Baker. He does not have permissions to create users.

4. Click **Logout**.
5. Log back in as the Web Console administrator.

Chapter 3: Deploy and run business flows

After your Web Console is properly configured, as described in [Configure the Web Console, on page 17](#), you are ready to begin deploying and running business flows. Business flows are a series of activities to be run in sequence. These activities include archive, copy, and reload cartridges, as well as other business logic in the form of Groovy scripts.

During and after the running of your business flows, you can monitor their status and review their history within the Web Console.

- [Before you begin](#)
- [Deploy business flows](#)
- [Run a business flow](#)
- [Run an undo business flow](#)
- [Run a reload business flow](#)

Before you begin

Before you can run a business flow, it must first be built in Designer along with the required cartridges and, if necessary, models and parameters. After you build the business flow, you need to deploy it in the environment where you plan to run the business flow. You can deploy a business flow from Designer or, if the deployment files are available, from the Web Console.

The tutorial in this section requires the following prerequisites:

- Before you can perform the steps in this tutorial, you must have configured the Web Console. See [Configure the Web Console, on page 17](#) for more information about configuring the Web Console.
- For this tutorial, two Oracle databases (distributed instance) are used. One contains the active data and the other the archive data. If you are using only one database (single instance) and/or another type of database, such as Sybase or SQL Server, you will notice some small differences in the dialogs but the basic principles are the same.
- This Web Console tutorial uses a business flow that was created in a Designer tutorial in the *HPE Structured Data Manager Developers Guide*. If you have not built that tutorial yourself yet, you need to do so before you proceed with this chapter.
- If you have not already done so, generate the deployment files for the archive, undo, and reload business flows from Designer:
- **Orders_D2D_BF**
- **Undo_Orders_D2D_BF**
- **Reload_Orders_D2D_BF**

Place the deployment files in a location that is accessible from the Web Console installation you are using. You will need this file path later on when you import the business flow in the Web Console. See the *HPE Structured Data Manager Developers Guide* for more information on how to generate deployment files.

Deploy business flows

To deploy a business flow

1. Launch the Web Console in your browser and login.
2. From the Web Console home page, click **Business Flow Management**.
3. If you have more than one environment in the Web Console, make sure that you have selected the correct one as the active environment. You can determine which environment is the active one and, if necessary, change it by clicking **ENVIRONMENT > Manage**.
4. Click **TASKS > Deployment > Deploy**.
5. If you see **Orders_D2D_BF**, **Undo_Orders_D2D_BF**, and **Reload_Orders_D2D_BF** in the list of business flows, select them. Otherwise, you must import these business flows and then select them. If necessary, click **Import** and follow the instructions in the online help to import the business flow into Web Console.

TIP: If the list of business flows is long, you can filter the list with the **Show latest versions only** and **Show newly imported business flows only**.

6. Click **Next**. The credentials page displays for you to enter your source and target database credentials, and define the access passwords.

For Oracle

« Back

Next »

Define passwords for Target and Archive Access Accounts

The passwords are used to provide access to Target and Archive Access schemas.

Target Access Password:*

Confirm Target Access Password:*

Archive Access Password:*

Confirm Archive Access Password:*

For SQL Server

« Back

Next »

Define passwords for Target and Archive Access Accounts

The passwords are used to provide access to Target and Archive Access schemas.

Target Access Password:*

Confirm Target Access Password:*

Archive Access Password:*

Confirm Archive Access Password:*

7. Enter the username and password for the source database administrator. These credentials are needed to access the active database.
8. Because a distributed instance (a separate archive database) is assumed for this tutorial, enter the username and password for the target database administrator. In a single instance environment

where the archive database and source database are one and the same, you do not need to perform this step because the source and target credentials are the same.

- 9. Create and confirm passwords for the target access user and the archive access user.
- 10. Click **Next**.
- 11. For Oracle, select the tablespaces for each business flow you are deploying.

For SQL Server or Sybase, enter the data and log size.

For Oracle

« Back

Next»

Define Tablespaces

Select Cartridges and choose tablespaces below

| Cartridges | History Data | History Temp | History Index | Archive Access Data | Archive Access Temp |
|------------------------------|--------------|--------------|---------------|---------------------|---------------------|
| <input type="checkbox"/> All | | | | | |
| Orders_D2D 1.0.0.0 | USERS (▼) | TEMP (▼) | USERS (▼) | USERS (73 MB) (▼) | TEMP (▼) |
| Orders_D2D_Reload 1.0.0.0 | USERS (▼) | TEMP (▼) | USERS (▼) | USERS (73 MB) (▼) | TEMP (▼) |

For SQL Server

« Back

Next»

Define Device Sizes

Select Cartridges and enter device sizes below

| Cartridges | History Data Size | History Log Size | Archive Data Size | Archive Log Size |
|------------------------------|-------------------|------------------|-------------------|------------------|
| <input type="checkbox"/> All | | | | |
| Orders_D2D 1.0.0.0 | 10 | 5 | 10 | 5 |
| Orders_D2D_Reload 1.0.0.0 | 10 | 5 | 10 | 5 |

- 12. Click **Next**. The Summary page displays for confirming the settings for this deployment.

For Oracle

«Back

Finish

Summary

Business Flows to be deployed:

Orders_D2D_BF 1.0.0.0

Contains Cartridges:

Orders_D2D 1.0.0.0

History Data Tablespace: USERS

History Temp Tablespace: TEMP

History Index Tablespace: USERS

Archive Access Data Tablespace: USERS

Archive Access Temp Tablespace: TEMP

Reload_Orders_D2D_BF_1.0.0.0

Contains Cartridges:

Orders_D2D_Reload 1.0.0.0

History Data Tablespace: USERS

History Temp Tablespace: TEMP

History Index Tablespace: USERS

Archive Access Data Tablespace: USERS

Archive Access Temp Tablespace: TEMP

Undo_Orders_D2D_BF 1.0.0.0

Contains Cartridges:

For SQL Server

«Back

Finish

Summary

Business Flows to be deployed:

Orders_D2D_BF 1.0.0.0

Contains Cartridges:

Orders_D2D 1.0.0.0

History Primary Data size: 10

History Transaction Log size: 5

Archive Access Primary Data size: 10

Archive Access Transaction Log size: 5

Reload_Orders_D2D_BF_ 1.0.0.0

Contains Cartridges:

Orders_D2D_Reload 1.0.0.0

History Primary Data size: 10

History Transaction Log size: 5

Archive Access Primary Data size: 10

Archive Access Transaction Log size: 5

Undo_Orders_D2D_BF 1.0.0.0

Contains Cartridges:

13. Click **Finish** to deploy the business flows. It may take some time for all three business flows to deploy.

Run a business flow

After you have deployed your business flows, you can run them.

To run a business flow

1. Click **TASKS > Launch > Business Flow**. The Business Flows page displays the list of deployed business flows, which you can run.

NOTE: The undo business flow, Undo_Orders_D2D_BF, is not displayed in this list. An undo business flow is associated with and can only be run in relationship to another business flow. For example, you can only run Undo_Orders_D2D_BF after Orders_D2D_BF has run. In [Run an undo business flow, on the next page](#), you will walk through the process of running an undo business flow.

2. Select the **Orders_D2D_BF** business flow. The Job Launch page displays for the selected business flow with its runtime parameter, Min_Months_to_Retain. The default value, 24, for the parameter is fine for this example.
3. Click **Run** to run the business flow.
4. Click **OK** to confirm.
5. Click **Monitoring** from the menu at the top of the page. The Business Flow Status page displays the business flow status and any currently running commands.
6. Click **Refresh** until the status of the business flow changes to COMPLETED. This status indicates that the business flow has completed successfully.

7. Query your archive database. The schema named DEMARC_HIST should contain tables populated with the archived rows from DEMARC.
8. Query your active database. The DEMARC tables should no longer contain the rows that were moved to DEMARC_HIST.

Run an undo business flow

Suppose that, after running your archive business flow and moving the data, you realize that you need to bring it all back into the active database. For example, perhaps you included too much data in the archive. In this circumstance, you want to quickly return the active and archive databases to the state they were in before you ran your archiving business flow. To return the moved data to the active database, you can run an undo business flow. You must create an undo business flow in Designer and deploy it to your environment. See the *HPE Structured Data Manager Developer's Guide*.

To run an undo business flow

1. Return to the **Web Console > Business Flow Management** page and click **Monitoring** in the top menu on the page.
2. Click the **Orders_D2D_BF** business flow to navigate to the Run Details page. On the Run Details page, you should see an **Undo** button above the business flow status table. This button only appears if an undo business flow is deployed.
3. Click **Undo**.
4. Click **OK** to confirm.
5. Click **Monitoring** from the menu at the top of the page. The Business Flow Status page displays the business flow status and any currently running commands.
6. Click **Refresh** until the status of the business flow changes to COMPLETED.
7. Query your archive database. The schema named DEMARC_HIST should have been returned to its state prior to having run the archive business flow. In this particular case, the tables should all be empty now.
8. Query your active database. The DEMARC schema should have been returned to its state prior to having run the Orders_D2D_BF business flow.

Run a reload business flow

Just because you archive a transaction today does not mean that it will not become active again at a later date. For example, you might have archived an order that later becomes active again due to a product recall or a customer payment issue. In such cases, you might need to bring the transaction back to the active database for further activity. To return archived data to the active database based upon some specified criteria, such as an order number, you can run a reload business flow. You must create a reload business flow in Designer and deploy it to your environment. See the *HPE Structured Data Manager Developer's Guide*.

TIP: It is always a best practice to create and deploy a reload business flow for each of your database to database archiving business flows. For more information about how to create a reload business flow, see the *HPE Structured Data Manager Developer's Guide*.

Before you begin

If you followed the steps in [Run an undo business flow, on the previous page](#), the target database no longer contains any data for you to reload. Before performing the steps in this section, you must re-run the Orders_D2D_BF business flow to populate the target database:

Follow the instructions in [Run a business flow, on page 46](#) to launch the archive business flow again. After it successfully completes, you can proceed with the rest of this tutorial.

To run a reload business flow

1. Return to the **Web Console > Business Flow Management** page.
2. Click **TASKS > Launch > Business Flow** from the menu at the top of the page. The Business Flows page displays the list of deployed business flows that you can run.
3. Select the **Reload_Orders_D2D_BF** business flow. The Job Launch page displays for the selected business flow with its runtime parameter, OrderID_to_reload.
4. For OrderID_to_reload, enter an ORDERID that you are certain was archived by the Orders_D2D_BF business flow, for example, 1402. If necessary, you can also query DEMARC_HIST in the archive database to find an ORDERID that was archived. Enter one of the ORDERID values that you find in the DEMARC_HIST.ORDER_HEADER table.
5. Click **Run**.
6. Click **Confirm**.
7. Click **TASKS > Monitoring > Overview**. The Business Flow Status page displays the business flow status and any currently running commands.
8. Click **Refresh** until the status of the Reload_Orders_D2D_BF business flow changes to COMPLETED.
9. Query the DEMARC_HIST.ORDER_HEADER table to ensure that the specified ORDERID has been removed.
10. Query the DEMARC.ORDER_HEADER table to ensure that the specified ORDERID has been restored.
11. If you plan to attempt other tutorials at this point, you need to perform one of the following procedures to restore the DEMARC sample data in the source database to its original state. Otherwise, you may get inconsistent results when you attempt the other tutorials.
 - Recreate the DEMARC schema by reloading the data with the load_demo script according to the instructions in Chapter 2, Configuring the Demarc data, in *HPE Structured Data Manager Tutorial*.
 - Undo the Orders_D2D_BF the same way that you did in [Run an undo business flow, on the previous page](#).

Chapter 4: Upload to HPE Vertica

Uploading data from your applications to HPE Vertica provides superior access and analytics. With HPE Structured Data Manager, you can set up jobs to take data from your database applications and quickly load the data into Vertica.

- [Introduction](#)
- [Create a Vertica location](#)
- [Deploy the business flow](#)
- [Specify a schema mapping file](#)
- [Run the business flow](#)

Introduction

To archive data to HPE Vertica, you just need an HPE Structured Data Manager instance installed and running on a system with access to HPE Vertica and the necessary privileges to write to the file system and insert data into the Vertica instance.

In the Web Console, you create a location that points to your Vertica database, which becomes the destination for the business flow.

TIP: When uploading to HPE Vertica, you should pre-create the schema into which you want to upload, for example `create schema demarc`.

Bear in mind the following considerations before uploading data into Vertica:

- Character columns (such as VARCHAR, TEXT, NTEXT, IMAGE, CLOB, and CHAR) cannot exceed 65000 bytes.
- Binary types (such as BLOB, BINARY, and VARBINARY) cannot exceed 12500 bytes.
- When running the same job to Vertica multiple times and inserting/modifying rows in existing tables, you need a data movement key for each table. Otherwise, only the first run will succeed and all of the subsequent runs will fail. In most cases, HPE Structured Data Manager can automatically select a unique key to use as the data movement key. In cases where it cannot, you may have to manually select one. For more information about data movement keys, see the chapter on working with models in the *HPE Structured Data Manager Developers Guide*.
- If your source database is Sybase and the tables you plan to archive make use of user-defined data types, you must ensure that those data types are available to the interface database created by HPE Structured Data Manager. For example, you might define the data type in the interface database or the model databases.


Create a Vertica location

To create a Vertica location to which you can upload data

1. Launch the Web Console in your browser and log in.
2. Create a new environment or make the environment from which you want to run your business flow the active environment, **Business Flow Management > ENVIRONMENT > Manage**.
3. Click **ENVIRONMENT > Active Environment > Locations**.
4. Click **New**.
5. Enter a name for the location, for example `salesorders_vertica_db`.
6. For **Type**, choose **Vertica**.
7. Enter the connection information for your Vertica database.
8. Click **Create**.

Deploy the business flow

To upload to HPE Vertica, you need a business flow with a database to file archive activity and an upload activity.



Deploy the business flow from Designer or Web Console. See the *HPE Structured Data Manager Developers Guide* or [Deploy business flows, on page 44](#).

Specify a schema mapping file

When uploading to HPE Vertica, you typically want to insert your data into a specific schema. For example, in our example, we want the data to go into a schema named `demarc` that you create in Vertica. If you do not specify a schema, the data will go into a general area of Vertica and you will likely have issues with subsequent upload operations.

Depending on your source database, you might also need to perform some mapping of the schema for the Vertica database. For example, suppose that your source data is in SQL Server and uses `DEMARC.dbo`, but, in Vertica, you simply want the schema to be `demarc`. A simple schema mapping file would take care of that difference.

To map your schema for upload to Vertica

1. Before running the upload to Vertica, open a vSQL session and create a schema called demarc:

```
salesorder=> create schema demarc;  
CREATE SCHEMA  
salesorder=> select schema_name from schemata;  
schema_name  
-----  
v_internal  
v_catalog  
v_monitor  
public  
demarc  
(5 rows)
```
2. Create a simple schema mapping file in a text editor called, for example, upload_vertica.properties with the following contents:

```
DEMARC.dbo=demarc
```
3. In the Web Console where you plan to run you database to file business flow and the upload to Vertica, go to **Business Flow Management > TASKS > System Parameters**.
4. Click **Cartridge**.
5. Click the cartridge from which you plan to upload data, for example, Orders_D2F.
6. Scroll down to **File Naming > Schema Mapping File for Upload** and enter the path and file you created in [step 2](#), for example:

```
C:\vertica_schema_mapping\upload_vertica.properties
```

| File Naming | | |
|--------------------------------|---|---------|
| Group File Prefix | <input type="text"/> | Default |
| Group File Suffix | <input type="text"/> | Default |
| Group XML Schema File Prefix | <input type="text"/> | Default |
| Group XML Schema File Suffix | <input type="text"/> | Default |
| Schema Mapping File for Upload | <input type="text" value="C:\vertica_schema_mapping\uplo"/> | Default |

7. Click **Apply**.

Run the business flow

To run your database to file business flow

1. Go to **Business Flow Management > TASKS > Launch Business Flow**.
2. Click the business flow you want to run.
3. Ensure that the **Destination DB Location** is properly set to the Vertica location you created

earlier.



| Runtime Parameters | |
|--------------------------|------------------------|
| Min_Months_to_Retain: | 72 |
| Destination DB Location: | salesorders_vertica_db |

4. Click **Run**. When prompted, confirm that you want to run the business flow. The business flow runs, archiving your data to file and then uploading it to Vertica.

Access the archived data

If the archive developer creates and deploys data access cartridges, users can use them to explore the archived data. Pre-defined data access is available by clicking **Data Access** from the home page in the Web Console. For information on how a developer can pre-create this kind of data access, see the *HPE Structured Data Manager Developers Guide*. Otherwise, you can simply query the data with vSQL to confirm that it was inserted into the database in the correct schema.

TIP: In order to run data access cartridges, users only require the Access Data privilege.

NOTE: To run a data access cartridge against Vertica, you must create a location for Vertica and the Vertica instance that you reference must be accessible from that instance of the Web Console.

For more information about HPE Vertica, see the following manuals:

- *HPE Vertica SQL Reference*
- *HPE Vertica Administrator's Guide*
- *HPE Vertica Programmer's Guide*

If you have a customer account, you can obtain these and other HPE Vertica documentation from my.vertica.com.

Chapter 5: Perform multiple data movements

Records from your database applications may be required for many different purposes, such as:

- **Compliance.** You may need to store a read-only copy of the data in a system like HPE Digital Safe or HPE Consolidated Archive for compliance purposes.
- **Data analytics.** You may want to upload the data to another database location, such as Vertica, in order to perform analytics on it.
- **Long term retention and query.** You may put infrequently accessed records into a file system where you can use the Archive Query Server to maintain SQL access to such records more cheaply than you would by keeping them in your database.

To meet these requirements, you can move data from a variety of sources to multiple targets in one job run using HPE Structured Data Manager.

- [Introduction](#)
- [Configure HPE Structured Data Manager](#)
- [Run business flows](#)

Introduction

A typical configuration that archives data to multiple remote locations consists of the following:

- A gateway instance of HPE Structured Data Manager running on a system with a data movement configuration that enables it to send data to various stores that you require, such as a network file system, HPE Digital Safe, HPE Consolidated Archive, and HPE Vertica.
- An application instance of HPE Structured Data Manager with one or more environments connected to your source, online transaction processing databases. One or more of the environments in this local instance include a location that points to the gateway HPE Structured Data Manager instance.

When the application instance targets jobs to the gateway instance, the gateway HPE Structured Data Manager picks up the data and performs the data movements specified in its configuration. This architecture enables you to satisfy multiple purposes from one archive run.

Configure HPE Structured Data Manager

To configure HPE Structured Data Manager for this case, you need two instances:

- [Configure a gateway instance of HPE Structured Data Manager](#)
- [Configure an application instance of HPE Structured Data Manager](#)

Configure a gateway instance of HPE Structured Data Manager

To configure your remote HPE Structured Data Manager

1. Install HPE Structured Data Manager as per the *HPE Structured Data Manager Installation Guide*.
2. Start the Web Console and install a repository as per the instructions in [Install the repository, on page 19](#).

TIP: Use PostgreSQL for the repository.

3. Create an environment for database to file as per the instructions in [Create an environment, on page 24](#).

TIP: Make the environment non-intrusive and use a JDBC URL to a PostgreSQL database for the source.

4. Create a Web Console user specifically to run the data movement jobs for the application instance of HPE Structured Data Manager, for example, `data_move_admin`. For more information, see [Create and edit Web Console users and groups \(optional\), on page 40](#).
5. Create locations to which you want to move data, for example HPE Autonomy Digital Safe, HPE Autonomy Consolidated Archive, and Vertica. See [Manage environments, on page 78](#) for more information.
6. Create a new directory in the home directory, under `/config/upload`, using the user name you created in [step 4](#), for example:

```
C:\ProgramData\Hewlett-Packard\HPOBT\config\upload\data_move_admin
```

7. Create a data movement configuration file in the directory you created in [Create a new directory in the home directory, under /config/upload, using the user name you created in step 4, for example, above](#). Use the sample file, `username.xml.sample`, which is found in `/config/upload/username`, as a template. Name the file for the target configuration you specify in [Enter the properties for the location, on page 56](#). For example, create a file called `sdm2sdm.xml`:

```
C:\ProgramData\HPE\HPEOBT\config\upload\data_move_admin\sdm2sdm.xml
```

Following is a simple example of the file contents:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
  <fileLocation environment="OLTP_Safe_Copy"
    location="LOCAL_ARCHIVE_FS" preserve="false"/>
  <dataMovements>
    <dataMovement destination="sdm_file_archive"/>
    <dataMovement destination="salesorders_vertica_db"/>
  </dataMovements>
</configuration>
```

The first element, `fileLocation`, tells HPE Structured Data Manager where it may place files for data movement for the specified environment, `OLTP_Safe_Copy`. `LOCAL_ARCHIVE_FS` is a file system

location defined in the Web Console and, by setting `preserve` to `false`, you let HPE Structured Data Manager know that it should remove the files after the data movement is done.

The `dataMovement` elements tell HPE Structured Data Manager where to move the archived data. Each `dataMovement` element references a location defined in the Web Console. For the purposes of this example, we created two such locations in [step 5](#):

- The first location is called `sdm_file_archive`, which happens to be a file system path.
- The second location is `salesorders_vertica_db`, a Vertica database. The archived data will be uploaded to this database.

NOTE: In the case of HPE Vertica, you should pre-create the schema to which you plan to upload the data. For example, suppose that the archived data came from an Oracle schema called DEMARC. You would create a DEMARC schema in the HPE Vertica database to hold the uploaded data before running the job.

These two examples do not require it, but some locations, such as HPE Digital Safe, may require a data transformation. Following is an example of a data transformation within a data movement:

```
<?xml version="1.0" encoding="UTF-8"?>
<configuration>
  <fileLocation environment="demarc"
    location="LOCAL_ARCHIVE_FS" preserve="false"/>
  <beforeDataMovements>
    println("Hello");
  </beforeDataMovements>
  <dataMovements>
    <dataMovement destination="LOCAL_TEMP_FS">
      <dataTransformation outputFormat="xml"
        batchSize="1">
      </dataTransformation>
    </dataMovement>
  </dataMovements>
  <afterDataMovements script="sayHello.groovy"/>
</configuration>
```

Notice also in this case the Groovy scripts to be run before and after the data movements. You can use Groovy to perform preparation and cleanup activities before and after you actual data movements.

Configure an application instance of HPE Structured Data Manager

To configure your local HPE Structured Data Manager

1. Install HPE Structured Data Manager as per the *HPE Structured Data Manager Installation Guide*.
2. Start the Web Console and install a repository according to the instructions in [Install the repository, on page 19](#).
3. Create an environment (standard or non-intrusive) for database to file as per the instructions in [Create an environment, on page 24](#).
4. For the environment you just created, go to the Environment page, **Business Flow Management**

> ENVIRONMENT > Active Environment.

5. Click **Locations** to create a location that points to your remote HPE Structured Data Manager.
6. Click **New**.
7. Enter the name to assign to the location, for example, `SDM_Safe`. The location name is case sensitive.
8. For **Type**, choose **SDM Gateway**.
9. Enter the properties for the location:
 - **Remote WebConsole URL** is the URL to the Web Console for your remote HPE Structured Data Manager. For example, `http://rovitto-linux.svs.usa.hp.com:8080/WebConsole`.
 - **Remote WebConsole user** is the user name on the gateway Web Console that you created in [step 4](#) to run the data movements, `data_move_admin`.
 - **Remote WebConsole password** is the password associated with the remote Web Console user.
 - **Proxy Server Host** is the host name of the proxy that must be used to reach a target HPE Structured Data Manager instance outside of your firewall.
 - **Proxy Server Port** is the port on the proxy server host that you are using to get outside of your firewall.
 - **Use SSL with client certificate** should be false, unless you have https configured for the Web Console. See [Configure HTTPS, on page 76](#) for information on how to set up https.
 - Optionally, enter **Target Configuration Name**, for example, `sdm2sdm`. This name corresponds to the configuration file name on the gateway HPE Structured Data Manager that you created in [step 7](#).
10. Click **Create**.

Run business flows

To launch your business flow on the application instance and send the data to the gateway instance

1. Deploy one or more database to file business flows in the local HPE Structured Data Manager, where the cartridges optionally reference an indexing cartridge. For more information on deployment, see the *HPE Structured Data Manager Developers Guide*.
2. Make sure the destination for the business flows is the location you created for the remote HPE Structured Data Manager.

TIP: To set the destination, go to **Business Flow Management > TASKS > System Parameters > Business Flow** and choose **Global BF Settings** or a specific business flow.

3. Run your business flows. See the *HPE Structured Data Manager Tutorial* or the Web Console help system for more information about running business flows.
4. Confirm that your archive has successfully run and the data pushed to all the specified locations in your data movement configuration file.

For example, in this case, check the file system for your gateway instance and query your HPE SDMVertica database.

Chapter 6: View eligibility analytics

During and after the running of your business flows, you can monitor their status and review their history within the Web Console. Furthermore, if enabled, you can examine any cartridge's eligibility analytics to see which records would be moved or copied based upon the selected rules.

This tutorial takes you through the process of viewing eligibility analytics, including:

- Deploying business flows
- Enabling eligibility analytics
- Running business flows
- Monitoring status
- Viewing eligibility analytics
- [Before you begin](#)
- [Enable eligibility analytics](#)
- [Run a business flow](#)
- [View eligibility analytics](#)
- [Continue the paused business flow](#)

Before you begin

Before you can run a business flow, it must first be built in Designer along with the required cartridges and, if necessary, models and parameters. After you build the business flow, you need to deploy it in the environment where you plan to run it. You can deploy a business flow from Designer or, if the deployment files are available, from the Web Console.

The tutorial in this section requires the following perquisites:

- Before you can perform the steps in this tutorial, you must have configured the Web Console. See [Configure the Web Console, on page 17](#) for more information about configuring the Web Console.
- This Web Console tutorial uses a business flow (**Orders_D2F_BF**) that was created in a Designer tutorial in Chapter 5, Eligibility analytics, of the *HPE Structured Data Manager Developers Guide*. If you have not built that tutorial yourself yet, you need to do so before you proceed with this tutorial.
- Make sure that the **Orders_D2F_BF** business flow is deployed from Designer (right click it and choose **Deploy**, then select the **Generate** radio button). Place the deployment file in a location that is accessible from the Web Console installation you are using.
- Deploy the **Orders_D2F_BF** business flow to the Web Console instance where you plan to run it, just as you would any other business flow.

Enable eligibility analytics

Using eligibility analytics can affect performance, so the default behavior at runtime is to not generate eligibility analytics. Therefore, if you want to generate the analytics, you must specifically change a parameter to enable them.

NOTE: You must also have enabled eligibility analytics in the rules in your model in Designer before deploying the business flow, and the model must contain a pause after the selection step in order for you to see them. For more information, see the eligibility analytics chapter in the *HPE Structured Data Manager Developers Guide*.

1. Launch the Web Console in your browser and log in.
2. From the Web Console home page, click **Business Flow Management**.
3. Click **TASKS > System Parameters**.
4. Click **Database to File**.
5. Scroll down to the section of the page entitled **Performance**.
6. For the Eligibility Analytics Configuration parameter, select **Enable the eligibility analytics**.

| Performance | | | |
|-------------------------------------|----------------------------------|---------|--|
| Combined delete count * | 100 | Default | |
| Combined statement count * | 50 | Default | |
| Data movement Batch size * | 50000 | Default | |
| Eligibility Analytics Configuration | Enable the eligibility analytics | Default | |

7. Click **Apply**.

Run a business flow

After you have enabled eligibility analytics, you can run the business flow.

1. Click **TASKS > Launch > Business Flow** from the menu at the top of the page.
The Business Flows page displays the list of deployed business flows.
2. Select the **Orders_D2F_BF** business flow.
The Job Launch page displays for the selected business flow with its runtime parameters, `Min_Months_to_Retain` and `Min_Selected_Rows`. The default values for the parameters are fine for this example.
3. Click **Definition** to view the Business Flow Definition page.
The Business Flow Definition page displays the activities that this business flow will run when you launch it. Notice the activity labeled **Suspend Job**. After records have been selected (but not moved) by the first activity, the business flow pauses, giving you the chance to review the eligibility analytics.
After you have reviewed the analytics and determined whether to continue, you can continue the business flow or cancel it.
4. Click **Back** to return to the launch page for the business flow.
5. Click **Run** to run the business flow.
6. Click **OK**.
7. Click **Monitoring** from the menu at the top of the page.

The Business Flow Status page displays the business flow status and any currently running commands.

8. Click **Refresh** until the status of the business flow changes to **SUSPENDED**. This status indicates that the business flow has paused as per the interrupt included when designing the business flow.

View eligibility analytics

With the business flow paused, you have the opportunity to review the eligibility analytics to determine whether you wish to continue running and move the data, or cancel the job without moving any data.

1. From the Business Flow Status page, click **Orders_D2F_BF** to navigate to the Run Details page. Notice how the Archive Selection Step and the Job Paused Step have completed, and the overall business flow is suspended.

2. From the Run Details page, click **Eligibility Analytics**.

The Eligibility Analytics page displays. Eligibility analytics are generated on a per-cartridge basis. Therefore, this page displays the cartridges in the business flow for which you can view eligibility analytics.

Notice the drop-down list next to the title of the table and next to each rule in the table.

3. Click the list next to the title of the table and select **Selection Data**.

The eligibility analytics for the cartridge displays.

4. If not already selected, select **All Data** in the left navigation pane.
 - Rows that were excluded by a rule appear in red and the name of the rule that caused their exclusion appears in the **Excluded By** column.
 - Rows that were selected for movement appear in black and the **Excluded By** column is empty.

TIP: To see only rows that were included (those that appear in black), select **Selection Only**. To see only rows that were excluded by the rules (those that appear in red), select **Exclusion Only**.

In some cases, you may find it useful to save the eligibility analytics for further review or future reference or both. You can easily save this data into an easily accessible format like XML or comma separated values (CSV).

5. Under **Export to**, click **CSV** to save the data to a CSV file.

CSV files can be opened or processed by many standard programs, such as spreadsheet applications. The operating system dialog for saving to a file displays.

6. Save the file under a name and location of your choice.
7. Open the CSV file in a spreadsheet program of your choice to further review the data.
8. (Optional) Explore the other options on the Eligibility Analytics page:
 - a. Return to the Run Details page by clicking **Monitoring > Orders_D2F_BF**.
 - b. Click **Eligibility Analytics**.
 - c. Click the various lists and select the other options to see what information is available to you.

Continue the paused business flow

For the purposes of this tutorial, assume that you are satisfied with the eligibility analytics review. You are now prepared to continue the business flow and actually move the data from the active database to the archive file.

1. Return to **Web Console > Business Flow Management > Monitoring > Overview**.
2. Click the **Orders_D2F_BF** business flow to navigate to the Run Details page.
3. Click **Continue** to complete the paused business flow.
4. Click **Monitoring > Overview** to return to the Business Flow Status page.
5. Click **Refresh** until the job finishes successfully.
6. Click the **Orders_D2F_BF** business flow to navigate to the Run Details page.
On the Run Details page, you can see the status of the individual activities within the business flow. In this case, all activities should have completed successfully and have a status of COMPLETED.
7. Click **Row Counts** to view the number of rows that were selected and deleted.

Run the undo business flow

In order to return the sample database to its original state, you should run the Undo_Orders_D2F_BF business flow.

1. Return to the **Web Console > Business Flow Management** and click **Monitoring** in the top menu on the page.
2. Click **OK**.
3. Click **Monitoring** from the menu at the top of the page.
The Business Flow Status page displays the business flow status and any currently running commands.
4. Click the **Orders_D2F_BF** business flow to navigate to the Run Details page.
On the Run Details page, the **Undo** button is displayed above the business flow status table.
5. Click **Undo**.

TIP: Alternatively, you could recreate the original DEMARC schema by reloading the data with the load_demo script according to the instructions in the *HPE Structured Data Manager Tutorial*.

Part III: Reference guide

This part provides an advanced task reference to assist you in performing specific tasks within the Web Console and from the command line. It includes:

- [Start and configure the Web Console](#)
- [Secure the Web Console](#)
- [Manage environments](#)
- [Manage business flows](#)
- [Run business flows and jobs from the command line](#)
- [Map schemas for upload or AQS cache](#)
- [Perform history schema reconciliation](#)
- [Use AQS to access database to file output](#)

Chapter 7: Start and configure the Web Console

The Web Console is the primary run time interface to HPE Structured Data Manager. You use it to deploy, run, and monitor jobs, and administer environments and users. The Web Console requires a repository in which to store its metadata. Hence, the first time you start the Web Console, you are prompted to install or upgrade the repository.

- [Start and stop the Web Console](#)
- [Install the repository](#)
- [Upgrade an existing repository](#)
- [Add Web Console nodes](#)
- [Use log files](#)

Start and stop the Web Console

Before you can connect to the interface for the Web Console server in your Web browser and perform tasks, you must start the Web Console server. You can also stop the Web Console server when it is not in use or when performing maintenance activities that require it be taken down.

Start the Web Console

You can start the Web Console from the Start menu on Microsoft Windows or from the command line on any platform.

In this section:

- [Before starting the Web Console](#)
- [Start Web Console from the program group](#)
- [Start Web Console from the command line](#)
- [Start Web Console as a Windows service](#)
- [Start Web Console on a different port, on page 65](#)

Before starting the Web Console

Before running the Web Console, ensure that you are logged into your operating system as a user with appropriate administrative privileges.

Start Web Console from the program group

1. After installing HPE Structured Data Manager, start the Web Console from:
Start > All Programs ><HPESDM_Pgm_Group> > Start Web Console

where `<HPESDM_Pgm_Group>` is the program group where you installed HPE Structured Data Manager.

A command window opens and closes automatically after the Web Console is started.

2. Launch a supported browser, and connect to the Web Console using the following default URL:

`http://<hostname>:8080/WebConsole`

where `<hostname>` is the name of the machine on which you installed HPE Structured Data Manager.

3. If you are starting Web Console for the first time after installation of HPE Structured Data Manager, you are prompted to either install a new repository or upgrade an existing one. See [Install the repository, on page 66](#) or [Upgrade an existing repository, on page 69](#).

Start Web Console from the command line

1. Open a command window and navigate to the directory that contains the Web Console script.

```
cd <install_dir>/obt/bin
```

where `<install_dir>` is the location where you installed the HPE Structured Data Manager.

2. Enter the start command:

On Windows:

```
webConsole.bat start
```

On UNIX:

```
./webConsole.sh start
```

HPE Structured Data Manager uses Tomcat to run the Web Console.

- On Windows, a new window is spawned for the Tomcat process.
- On Linux, a Tomcat process is started.

Check the Tomcat process and logs in your installation directory to determine if the Web Console started successfully.

3. Launch a supported browser, and connect to the Web Console using the following default URL:

`http://<hostname>:8080/WebConsole`

where `<hostname>` is the name of the machine on which you installed HPE Structured Data Manager.

4. If you are starting the Web Console for the first time after installation of HPE Structured Data Manager, you are prompted to either install a new repository or upgrade an existing one. See [Install the repository, on page 66](#) or [Upgrade an existing repository, on page 69](#).

Start Web Console as a Windows service

1. On Windows, install the Web Console service, if it is not already installed.

```
<install_dir>\obt\bin\webConsoleService.bat install
```

TIP: To uninstall the service, run `webConsoleService.bat uninstall`.

2. Go to **Control Panel > Administrative Tools > Services**.
3. Find the Web Console service, HPE Structured Data Manager. By default, the service is

configured to start automatically during system start. You can also start it manually from the Control Panel or from the command line using `webConsoleService.bat start`.

4. Launch a supported browser, and connect to the Web Console using the following default URL:

`http://<hostname>:8080/WebConsole`

where `<hostname>` is the name of the machine on which you installed HPE Structured Data Manager.

5. If you are starting the Web Console for the first time after installation of HPE Structured Data Manager, you are prompted to either install a new repository or upgrade an existing one. See [Install the repository, on the next page](#) or [Upgrade an existing repository, on page 69](#).

Start Web Console on a different port

By default, the Web Console runs on port 8080. If you want to use a different port number, you must change it in the Tomcat configuration file.

To change the port number

1. Update the ports in `server.xml` in `<install_dir>\obt\webconsole\apache-tomcat\conf`.
2. Start the Web Console as you normally would, for example, from the command line using `webConsole.bat/sh`.

3. Launch a supported browser, and connect to Web Console using the following URL:

`http://<hostname>:<port>/WebConsole`

where `<hostname>` is the name of the machine on which you installed HPE Structured Data Manager, and `<port>` is the port number you specified in the previous step.

4. If you are starting the Web Console for the first time after installation of HPE Structured Data Manager, you are prompted to either install a new repository or upgrade an existing one. See [Install the repository, on the next page](#) or [Upgrade an existing repository, on page 69](#).

Stop the Web Console

You can stop the Web Console from the Start menu on Windows or from the command line on any platform.

Stop the Web Console from the Start menu

1. Stop Web Console from the Start menu:

Start > All Programs ><HPESDM_Pgm_Group> > Stop Web Console

where `<HPESDM_Pgm_Group>` is the program group where installed HPE SDM.

A command window opens and closes automatically after the Web Console is started.

Stop the Web Console from the command line

To stop the Web Console, use the following command:

`webConsole.bat/sh stop`

NOTE:

Before installing the repository, you can set the password policies in `<SDM_dir_HOME>/config/WebConsole.properties` file. For more information, see [Set password security](#)

Install the repository

After you install HPE SDM, and start and launch the Web Console, you are prompted to install a new repository or upgrade an existing one. For information about upgrading an existing repository, see [Upgrade an existing repository, on page 69](#).

You can create the repository on an Oracle, SQL Server, or PostgreSQL database. You have the option to connect with the standard Oracle, SQL Server, or PostgreSQL connection formats, or to use a JDBC URL.

1. Launch the Web Console in your browser.
2. Click **Install a New Repository**.
The Repository Database: Administrator page displays.
3. Select the RDBMS type for your repository and enter the appropriate information:
 - [Oracle, below](#)
 - [SQL Server, on the next page](#)
 - [PostgreSQL, on page 68](#)
 - [JDBC URL, on page 68](#)

Oracle

- a. Enter the following administrator information:

| Field | Description |
|--------------|--|
| User | Enter the administrator username. |
| Password | Enter the password for the specified user. |
| Host | Enter the name of the machine where your Oracle database is installed. |
| Port | Enter the port number of your Oracle database, for example, 1521. |
| Service Name | Enter the name of your Oracle database, for example, ORCL. |

- b. Click **Next**.
The Repository Database: User page displays.
- c. Select or enter the following information for the repository user:

| Field | Description |
|-------|--|
| User | Enter the name you want to use for the repository user. The default value is <code>obt_rep</code> . |

| Field | Description |
|------------------------|--|
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Data Tablespace | Select the data tablespace you want to use. |
| Temp Tablespace | Select the temporary tablespace you want to use. |
| Encryption Key | Enter the encryption key you want to use. |
| Confirm Encryption Key | Enter the encryption key again. |

SQL Server

- a. Enter or select the following administrator information:

| Field | Description |
|---------------|---|
| User | Enter the administrator username. Required for SQL Server authentication only. |
| Password | Enter the administrator password for the specified user. Required for SQL Server authentication only. |
| Host | Enter the host name or address of the machine where your SQL Server database is installed. |
| Port | Enter the port number of your SQL Server database. |
| Instance Name | Enter the name of your SQL Server. |

- b. Click **Next**.

The Repository Database: User page displays.

- c. Select or enter the following information for the repository user:

| Field | Description |
|------------------------|---|
| User | Enter the name you want to use for the repository user. The default value is obt_rep. |
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Primary Data Size | Accept the default value or enter a maximum size in MB for the data file. |
| Transaction Log Size | Accept the default value or enter a maximum size in MB for the log file. |
| Encryption Key | Enter the encryption key you want to use. |
| Confirm Encryption Key | Enter the encryption key again. |

PostgreSQL

- a. Enter or select the following administrator information:

| | |
|---------------|--|
| User | Enter the database administrator username. |
| Password | Enter the database administrator password. |
| Host | Enter the name or address of the machine where the database is installed. |
| Port | Enter the port number of your PostgreSQL database. |
| Database Name | Enter the name of the PostgreSQL database, for example, <code>sdm_repos</code> . You can only have one repository per PostgreSQL database. |

- b. Click **Next**.

The Repository Database: User page displays.

- c. Select or enter the following information for the repository user:

| Field | Description |
|------------------------|--|
| User | Enter the name you want to use for the repository user. The default value is <code>obtrep</code> . |
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Encryption Key | Enter the encryption key you want to use. |
| Confirm Encryption Key | Enter the encryption key again. |

JDBC URL

- a. Enter the following administrator information:

| Field | Description |
|----------|--|
| User | Enter the administrator username. |
| Password | Enter the administrator password for the specified user. |
| URL | Enter the URL for the database. For more information, see Use JDBC URLs for database connections, on page 80 . |

- b. Click **Next**.

The Repository Database: User page displays.

- c. Select or enter the following information for the repository user:

| Field | Description |
|-------|---|
| User | Enter the name you want to use for the repository user. The default value is <code>obt_rep</code> . |

| Field | Description |
|------------------------|--|
| Password | Enter the password for the user you created. |
| Confirm Password | Enter the password again. |
| Data Tablespace | Select the data tablespace you want to use. |
| Temp Tablespace | Select the temporary tablespace you want to use. |
| Encryption Key | Enter the encryption key you want to use. |
| Confirm Encryption Key | Enter the encryption key again. |

4. Click **Next**.

The Console Administrator page displays.

5. Enter the following information:

| Field | Description |
|------------------|--|
| User Name | Enter the name you want to use for the Web Console administrator. The Web Console administrator is used to log on to the Web Console. |
| Password | Enter the password for the administrator. |
| Confirm Password | Enter the password again. |
| Real Name | Optionally, enter the name of the administrator user. |
| Description | Optionally, enter a description. |
| Email | Optionally, enter an email address. |
| Phone Number | Optionally, enter a phone number. |

6. Click **Next**.

The Summary page displays.

7. Click **Finish** to create the repository, or click **Back** to make changes.

The following message displays:

You have successfully installed the repository. Please do not close your browser. Web Console will restart.

After the Web Console restarts, the login screen displays.

Upgrade an existing repository

The first time that you start the Web Console after installing HPE Structured Data Manager, you must install a new repository or import one from an earlier version. For information on installing a new repository, see [Install the repository, on page 19](#).

For information on upgrading an existing repository, see the *HPE Structured Data Manager Upgrade Guide*.

Add Web Console nodes

In some cases, you may want to have additional Web Console nodes on other machines that operate upon the same repository. Having multiple Web Console nodes enables you to perform operations against the same repository from any one of several different machines, thus avoiding any performance bottlenecks on the primary Web Console node. By copying certain files from your first or primary node of the Web Console, you can effectively clone the primary node on other machines.

NOTE: This procedure assumes that you have already installed and fully configured one instance of HPE Structured Data Manager according to the instructions in the *HPE Structured Data Manager Installation Guide* and [Install the repository, on page 66](#).

To configure an additional Web Console node

1. From the primary node, where you first installed HPE Structured Data Manager and configured the Web Console, make copies of the following files in the HPE Structured Data Manager program data directory and move them to a location accessible to the machine where you want to create an additional node:
 - <prime_app_data>\HPOBT\config\connection-sources.xml
 - <prime_app_data>\HPOBT\config\hpdbackiving.auth
 - <prime_app_data>\HPOBT\config\sql_tuning.properties
2. On the machine where you want to create the additional Web Console node, use the Installer to create and populate an installation directory tree just as you would for a new installation. Follow the instructions in the *HPE Structured Data Manager Installation Guide*, but do not start the Web Console.
3. Apply all of the patches that were applied to the original code tree.

NOTE: Be aware that the patches will attempt to copy new files into the code tree and that some patches will then attempt to execute scripts to perform tasks (such as compiling packages in databases). However, these scripts will fail when executed on servers that are strictly Web Console servers; this is not a problem, though, because the critical task that must be completed is simply that the files are copied.

4. Overwrite the following files on the clone machine with the copies you made from the primary node:
 - <clone_app_data>\HPOBT\config\connection-sources.xml
 - <clone_app_data>\HPOBT\config\hpdbackiving.auth
 - <clone_app_data>\HPOBT\config\sql_tuning.properties
5. Start the cloned Web Console node as described in [Start and stop the Web Console, on page 63](#).
6. Confirm that you are operating upon the same repository as the primary node. You should not be prompted to install or upgrade a repository as you normally would for a new installation of the Web Console. Furthermore, if you are logged in as the administrator user, you should be able to see any environments that were created from the primary node.

Use log files

When reviewing job status or diagnosing an issue, it can sometimes be helpful to view the complete information provided in log files for HPE Structured Data Manager. If you are regularly consulting the log files for a particular purpose, you may also want to adjust the logging properties to include or exclude certain information.

View log files

Error messages from the Web Console and HPE Structured Data Manager are captured and appended to log files. By default, the log files are saved to the home directory under `log`. See [Manage the home directory, on page 39](#).

To view log files

1. Navigate to the directory containing the log file.

Example

```
cd <app_data_dir>/HPE/HPEOBT/log/  
or  
cd <app_data_dir>/HPE/HPEOBT/log/<envName>/
```

2. Open one of the following log files using a text editor of your choice:

| File name | Description |
|--|--|
| obt.log | Captures logging information that is not specific to an environment. For example, repository install or Web Console. New lines appended to the log file are displayed as they appear. Multiple log files are numbered in sequence, for example, obt.log1. |
| env_ <envName>.log | Captures logging information that is specific to an environment but not to a run. For example, deploying a business flow. The environment log file is located in the home directory under <code>log/<envName></code> . See Manage the home directory, on page 39 . |
| <envName>_run_ <runID>.log | Captures logging information for a specific business flow run. The business flow log file is located in the home directory under <code>log/<envName></code> . |
| pdm_server_ <runID>.log where <runID> is the run ID of the job as displayed on the Web Console. | The <code>pdm_server_<runID>.log</code> file is used for advanced data selection. The log file contains the generated statements, execution plans, and statistics from the execution of the advanced data selection. |

Edit the obt.log logging properties

You can edit the `log4j.properties` file to change the following logging properties for the `obt.log` file:

- Where the log files are kept
- What information is logged
- Maximum size of the log file
- How much logging information is kept

By default, the `obt.log` file is limited to 10 MB. When it exceeds the default limit, the `obt.log` file is renamed to a backup file, and a new `obt.log` file is created. By default, a total of nine backup files are kept, limiting the log files to a 100 MB maximum size.

To edit the obt.log properties

1. Navigate to `config` under the home directory, which contains the `log4j.properties` file. See [Manage the home directory, on page 39](#).
2. Create a backup copy of `log4j.properties` that you can roll back to in case of any problems.
3. Open the `log4j.properties` file with a text editor of your choice.
4. Make your desired changes to the file.
5. Save the `log4j.properties` file. The changes are applied automatically.

You can edit the `outerbay.properties` file to control the number of log files.

To set the number of log files

1. Navigate to `config` under the home directory, which contains the `outerbay.properties` file. See [Manage the home directory, on page 39](#).
2. Create a backup copy of `outerbay.properties` that you can roll back to in case of any problems.
3. Open the `outerbay.properties` file with a text editor of your choice.
4. Set the `envBasedLogging` and `runBasedLogging` properties to control the number of logging files as follows:
 - To have only the log file `obt.log`, set both `envBasedLogging` and `runBasedLogging` to false.
 - To have the log file `<envName>.log`, set `envBasedLogging` to true.
 - To have the log file `<envName>_run_<runID>.log`, set `runBasedLogging` to true.
 - To have all three logs, set both `envBasedLogging` and `runBasedLogging` to true.
5. Save the `outerbay.properties` file. The changes are applied automatically.

Chapter 8: Secure the Web Console

To more securely and conveniently manage access to and communication with the Web Console, you can implement LDAP and configure HTTPS. You can also enforce stricter password controls.

- [Import a certificate](#)
- [Implement LDAP](#)
- [Set password security](#)
- [Configure HTTPS](#)
- [Configure HPE IDOL security](#)
- [Restrict Web services for cloud deployments](#)

Import a certificate

To import a certificate for the Web Console to use:

1. Obtain a certificate from a valid authority and copy to a file in *install_dir/jre/lib/security*.
2. Open a shell and set JAVA_HOME to HPE Structured Data Manager's JRE. For example, on Windows:

```
set JAVA_HOME="c:\Program Files\HPESDM740\jre\bin"
```

3. Ensure that this same JRE path is also appended to the PATH in the shell. For example, on Windows:

```
set path=%PATH%;C:\Program Files\HPESDM720\jre\bin
```

4. Change directories to *install_dir/jre/lib/security*, where you placed the certificate.
5. Run keytool again to import the certificate:

```
..\..\keytool -importcert -keystore cacerts -file <certfile>
```

NOTE: For more information about keytool, see your Java documentation.

Implement LDAP

LDAP allows you to leverage your existing corporate directory to grant users access to the Web Console. To enable LDAP for use with the Web Console, you need to perform the following tasks:

- [Import a certificate](#)
- [Update webConsole.properties](#)
- [Add users](#)

NOTE: The instructions that follow assume the use of LDAPS because it is best practice for security, but you can implement LDAP without SSL if your environment allows it and security is not a major concern.

Update webConsole.properties

Now that you have done all of the preparation, you are ready to cut over to LDAP or LDAPS for access control:

1. Stop the Web Console. See [Start and configure the Web Console, on page 63](#).
2. In a text editor, open `webConsole.properties`, which is stored in the application data directory.
3. Uncomment the LDAP configuration and update the values for your environment. Note that `security.ldapServer` can be `ldap` or `ldaps`. For example:

NOTE: The LDAP settings that you need to include may vary depending upon your LDAP setup. Consult with your LDAP administrator to ensure you have the right settings specified in `webConsole.properties`.

```
security.ldapRetrieveDatabaseRoles=true
security.ldapRetrieveGroupRoles=true
security.ldapServer=ldaps://MYldap.my.com:636
security.ldapManagerDn=
security.ldapManagerPassword=
security.ldapSearchBase=ou=People,o=my.com
security.ldapSearchFilter=(uid={0})
security.ldapGroupSearchBase=ou=Groups,o=my.com
security.ldapGroupSearchFilter=(member={0})
security.ldapPasswordAttributeName=<as appropriate>
```

4. Start the Web Console. See [Start and configure the Web Console, on page 63](#).

Include the LDAP manager password

1. If you want to perform authenticated LDAP searches, you need to enter your LDAP manager password in `webConsole.properties` as follows:
`security.ldapManagerPassword=password`
2. With `ldapManagerPassword` set, administrators are prompted upon their next login to choose whether to encrypt the password in the `webConsole.properties` file.

NOTE: If you have multiple administrators, each one is prompted upon login.

- If an administrator chooses **Yes**, the `ldapManagerPassword` setting with the password in clear text is removed from `webConsole.properties` and the `ldapEncryptedManagerPassword` is appended to the end of the file:

```
security.ldapEncryptedManagerPassword=<encrypted value>
```

CAUTION:

Encryption of the LDAP manager password in `webConsole.properties` is not intended as a substitute for file security. All of your HPE Structured Data Manager configuration files, including `webConsole.properties`, should be stored in secured directories with restricted access.

- If an administrator chooses **No**, the `ldapManagerPassword` setting with the password in clear text remains intact and the administrator will receive no further prompts to encrypt it.

Add users

After you have enabled LDAP and restarted the Web Console, you can begin to add Web Console users and groups from your corporate directory.

1. Navigate to **Business Flow Management** in the Web Console.
2. Click **Users**.
3. Add users as described in the Web Console online help. Ensure that you are using names and domains that exist in your LDAP directory, for example, john.smith@yourdomain.com.

NOTE: You should maintain at least one built-in user that you can use to log in to the Web Console should your LDAP directory go down for any reason.

Set password security

To further increase security, you can place some additional requirements on your login and password behavior.

To set password security

1. Stop the Web Console. See [Start and configure the Web Console, on page 63](#).
2. In a text editor, open `webConsole.properties`, which is stored in the home directory under `config`. See [Manage the home directory, on page 39](#).
3. Modify the following properties as desired to control login and password behavior:

Table 1 Login and password properties

| Property | Description |
|--|---|
| <code>security.user.authMaxFailedAttempts</code> | Defines the maximum number of failed login attempts allowed before the user's account is locked. The default value is 5. |
| <code>security.password.minLength</code> | Defines the minimum required length of a password. The default is 8. |
| <code>security.password.mustHaveLetter</code> | Indicates that the password must have at least one letter. |
| <code>security.password.mustHaveDigit</code> | Indicates that the password must have at least one digit. |
| <code>security.password.ruleMessage</code> | Is the message displayed to users when their new password is rejected for some reason. Typically this message outlines the password requirements. |

Table 1 Login and password properties, continued

| Property | Description |
|---------------------------------------|---|
| security.password.mustHaveSpecialChar | Indicates that the password must have at least one special character out of !@#\$%&* _+=?~- |

NOTE:

Security policies can be added even before installation of repository, but if the HOME folder is changed during the installation of repository, then you must set the password security policies again in the `webConsole.properties` file which is stored in the `HOME/config` directory.

5. Start the Web Console. See [Start and configure the Web Console, on page 63](#).

Configure HTTPS

The Web Console runs on an instance of Tomcat, which is located in `<install_dir>\obt\webconsole\apache-tomcat`. You implement SSL just as you would for any other Tomcat instance. See Tomcat documentation and the Tomcat `server.xml` configuration file.

IMPORTANT:

When configuring Tomcat for SSL, you should avoid using the APR implementation because it uses the OpenSSL engine by default.

Configure HPE IDOL security

If you use HPE Structured Data Manager to index structured records for IDOL, you should check your IDOL settings in Web Console and choose appropriate security options for communicating with IDOL.

To review and change your IDOL settings for the Web Console

1. Navigate to **Business Flow Management** in the Web Console.
2. Click **Settings** from the menu at the top of the page.
3. Click **IDOL Settings** from the left navigation pane. The IDOL Settings page displays.
4. Select the protocol to employ when communicating with the IDOL server, HTTP or HTTPS.
For more information about HTTPS and IDOL See the *HPE IDOL Administration Guide*, "Set Up an SSL Connection".
5. If you have configured TEA encryption for IDOL, you can enable it here by checking **Enabled** and entering the key in the provided field.

TEA encryption is configured in `AutonomyIDOLServer.cfg`. For example:

```
ACIEncryption]
CommsEncryptionTEAKeys=1111,2222,3333,4444
CommsAllowUnencrypted=false
CommsEncryptionType=TEA
```

6. Optionally, you can change the Security by Database setting as well:
 - **Allow users with data access privileges to query any IDOL database** means that any user who has the data access privilege granted to them can also query any IDOL database. See the Web Console online help for more information on privileges.
 - **Grant IDOL database query access** means that you must explicitly grant users query access to IDOL databases. If you choose this option, when you select **Apply**, a new link appears called **Set security by database**. Following the link takes you to a page where you can manage query access by IDOL database and user.
7. Click **Apply**.

Restrict Web services for cloud deployments

When deployed in the cloud, you should disable deployments and job runs from the Web Console:

1. Make the following changes in `outerbay.properties`:

```
allow.remoteDeployment=false  
allow.remoteRun=false
```

By changing these values to false, the server will refuse to do deployments and start business flows using Web services.

Chapter 9: Manage environments

An environment is a named deployment environment associated with a source (active) database. If you plan to perform database to database operations for the source database, the environment will also include a target database or location. From an environment, users with the necessary privileges can run business flows against the source database. You can create as many environments as you wish, thus supporting multiple source and target databases from a single installation of HPE Structured Data Manager.

Before you can run business flows, you must create at least one environment into which you can deploy the business flows.

TIP: Step-by-step information for the Web Console user interface can be found in the Web Console help system.

- [Choose environment type](#)
- [Use JDBC URLs for database connections](#)
- [Configure standard database to file environments for DB2](#)
- [Configure parameters](#)
- [Manage locations](#)
- [Refresh test environments](#)

Choose environment type

In Web Console, you can have multiple, different combinations of data source and target. For example, you might have ten source databases from which to archive and five locations in which to store the archived data. Environments provide the mechanism for you to uniquely identify a specific combination of source and its potential targets. When you run jobs, you run them in a particular environment, that is, against a specific combination of source and target. Database to file archiving is enabled for each environment and you can optionally add database to database archiving to the environment.

Before you start, you must consider what type of environment best meets your business requirements and policies. HPE Structured Data Manager offers two major types of environment:

Standard environments

In a standard environment, data is archived using an interface schema created in your source database, and native or JDBC drivers to move the data. Standard environments support all cartridge types, including database to database cartridges.

When to use

A standard environment is best when:

- The data source has no restrictions on the operations that you can perform, such as DELETE.
- You require database to database archiving.

- Performance is critical.

Non-intrusive environments

In a non-intrusive environment, data is archived without creating an interface schema in the source database. Non-intrusive environments use the corresponding JDBC driver for the source database. You can use the JDBC driver that comes with HPE Structured Data Manager for natively supported platforms. If you are using a non-native database, you must supply your own JDBC driver.

NOTE: Non-intrusive environments only support database to file archiving. You cannot run database to database cartridges in a non-intrusive environment.

Non-intrusive environments support all of the following functionality:

- Data masking
- Uploading of extracted files to supported databases.
- Use of Column Inclusion, Name Override and Custom Properties.
- Delete operations, provided it is possible to perform delete operations against the data source. If you attempt to delete something when your source does not support deletes, you will receive an error message when running the delete operation. Delete-only cartridges are supported as well, but they are subject to these same limitations.

When to use

A non-intrusive environment is best when:

- Database to database archiving is not required.
- You must copy from read-only sources, especially in cases where data is associated with older technologies that might not support basic SQL statements such as DELETE, or when the database administrator or company policy prohibits creating objects in the production database.
- Performance is not a major issue. Standard environments may perform better than non-intrusive environments in some instances. If performance is a problem and the data source is not read-only, you may want to consider a standard environment.
- None of the following behaviors presents a problem for you:
- Selection and data movement occur simultaneously. If you split an archive activity into separate selection and data movement segments, selection is automatically skipped.
- When deploying a database to file cartridge against a non-intrusive environment, no selection tables are created. Data movement is performed directly from the source environment.
- Because there is no interface schema, Eligibility Analytics are not supported. Also, Validation is not available. If you specify Validation in a database to file cartridge, it will be ignored when the cartridge is deployed.
- Reload and undo are not supported. However, you can upload extracted XML or CSV files. After extracting data from a non-intrusive environment, you can upload the files to any of the supported database platforms.
- Advanced Selection and chaining tables are not supported. If you specify Advanced Selection, cartridge deployment will fail.
- Custom selection is supported in non-intrusive environments; however, you must include a SELECT

statement against the driving table (rather than an INSERT statement). The SELECT statement must include all columns from the driving table. If all columns are not selected in the driving table, a runtime error occurs when you attempt to deploy the cartridge. The error occurs because the custom SQL is not parsed, and it cannot be determined whether all columns are included at the time of deployment.

Use JDBC URLs for database connections

You can use a JDBC URL to connect to your repository and databases. For your repository, the connection can only be to Oracle, SQL Server, or PostgreSQL.

JDBC URLs

To use a JDBC URL for your connections, you need to enter the following information in the appropriate format:

NOTE: You can only use one JDBC driver at a time. Therefore, the URL you use to invoke the driver must be unique.

| Database | Expected URL |
|------------|---|
| Oracle | <code>jdbc:oracle:thin:@<Host>:<Port>:<OracleServiceName></code> |
| SQL Server | <code>jdbc:sqlserver://<Host>:<Port>;InstanceName=<DatabaseName></code> or <code>jdbc:hpdbsqlserver://<Host>:<Port>;InstanceName=<DatabaseName></code> <div>NOTE:<ul style="list-style-type: none">Copy the SQLJDBC4.JAR driver from <code>/obt/lib</code> to <code>/obt/lib/jdbc</code>.The <code>sqlserverToPostgresql.xml</code> file from <code>..\obt\foundation\components\config\AQSCache</code> location needs to be copied and named as <code>ansiToPostgresql.xml</code> if the source database is <code>sqlserver</code> and target database is <code>Postgresql</code>.</div> |
| Sybase | <code>jdbc:sybase://<Host>:<Port>;InstanceName=<DatabaseName></code> |
| DB2 | <code>jdbc:db2://<Host>:<Port>/<DatabaseName></code> |

NOTE: For repository creation, you can only use Oracle or SQL Server.

| Parameter | Description |
|-------------------|--|
| Host | The fully qualified host name or address for the connection. |
| Port | The port number for the connection. |
| OracleServiceName | The Oracle database service name. |
| DatabaseName | The name of the SQL Server, Sybase, or DB2 database. |

Oracle example

```
jdbc:oracle:thin:@localhost:1521:ORCL
```

Oracle RAC

You can use database to database archiving in an Oracle RAC environment that has more than one server, or node, connected to a database on a shared file system. To ensure connectivity even in the event of a rollover:

- The repository must be installed either on a non-RAC server or using a JDBC URL connection.
- The active database must use a JDBC URL connection.
- For distributed archive, the archive database must use a JDBC URL connection.

Contact your database administrator for the correct JDBC URL that you should use. The JDBC URL for Oracle RAC should look similar to this:

Example

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=Node1)
(PORT=1521))(ADDRESS=(PROTOCOL=TCP)(HOST=Node2)(PORT=<port>))(LOAD_BALANCE=no)
(FAILOVER=yes))(CONNECT_DATA=(SERVICE_NAME=My_Database)))
```

Configure standard database to file environments for DB2

When you create a standard database to file environment with DB2 as the source, you must set two additional properties in <install_dir>\obt\install\product_<user>.properties:

- source.interface.default.schema
- source.interface.default.storage.location

Properties file example

The following is an example of product_<user>.properties:

```
#Product Information
encryption.key = hello
environment.id = DefaultEnvironment
environment.description = my db2 environment
ea.product.configuration = default
##ea.product.configuration = lite

#Source Database Information
source.rdbms.name = db2
##source.rdbms.name = db2

source.conn.dbserver = db2i97a
source.conn.host = linux-qacert-08
source.conn.port = 50000
#source.conn.conntype=url
```

```
#source.conn.url = jdbc:db2://linux-qacert-08:50000/db2i97a
source.dbadmin.username = db2i97a
source.dbadmin.password = all4one

#Interface User Information
#source.interface.dbname = db2i97a_SOURCE_INTF
source.interface.default.schema=${source.interface.default.schema}
source.interface.default.storageLocation=${source.interface.default.storageLocation}
n}
##source.interface.dbname =
source.interface.owner.name = db2rt
source.interface.owner.password = welcome
##source.interface.owner.name = db2i97a
##source.interface.owner.password = all4one
```

Configure parameters

Before you run business flows, you should review and configure the parameters that govern the running of those business flows. HPE SDM provides parameters at the following levels:

- At the database to database level, you configure parameters that apply to all database to database business flows and cartridges. For a listing of all of these parameters, see [Configuration and runtime parameters, on page 139](#).
- At the database to file level, you configure parameters that apply to all database to file business flows and cartridges. For a listing of all of these parameters, see [Configuration and runtime parameters, on page 139](#).
- At the business flow level, you configure the user-defined configuration parameters created in Designer, and those parameters apply to all business flows. For more information, see the *HPE Structured Data Manager Developer's Guide*.
- At the cartridge level, you configure parameters that apply to the particular instance of the cartridge, and override the value of the database to database level or database to file level parameter. If the same cartridge belongs to more than one business flow, each instance of the cartridge can have different parameter values.

Choose a data movement method for database to database

One database to database parameter change that you might need to make for performance reasons is the Data movement method parameter in the Performance section. HPE SDM supports the following data movement options for database to database:

- fully transactional
- table parallel

When using the fully transactional method, all data, including LONGs, are copied and deleted in the same transaction. The data is copied directly from the OLTP tables to the history tables, and the history data is always consistent.

The fully transactional method is best when consistent access to the archived data must be maintained.

When using the table parallel method, different tables are copied in different transactions in parallel, directly from the active database to the archive. Unlike the fully transactional method, the data is deleted in a separate transaction. Because the copy and delete are in separate transactions, there may temporarily be missing or duplicate data in the history tables before the entire process is complete.

NOTE: When using the table parallel method, the transparency layer should not be accessed until the data movement process is complete, as it may display errors or show inconsistent data.

The table parallel method is best when faster performance is required, and temporarily inconsistent historical data is not an issue.

To switch between data movement methods

1. Access the Parameters page as described in [Choose a data movement method for database to database, on the previous page](#).
2. Click **Database to Database** from the left navigation pane.

The Parameters - Database to Database page displays.

3. Change the value of the Data movement method parameter in the Performance section.
4. Click **Apply**. The configuration parameters are saved.

Configure user-defined business flow parameters

In Designer, you can create configuration parameters that are used as part of a business flow that is deployed into a particular environment. These parameters appear on the business flow level of the Parameters - Business Flows page. If no user-defined configuration parameters were created for your business flows, this page is empty.

For more information on user-defined parameters, see the *HPE Structured Data Manager Developers Guide*, or contact your archive developer.

Configure cartridge parameters

The cartridge level parameters apply to one instance of the selected cartridge, and are displayed on the Cartridge Configuration page. Except for the file naming parameters, all cartridge level parameters are a subset of the component level parameters listed in [Configuration and runtime parameters, on page 139](#).

Because file naming parameters create the naming convention for each individual database to file cartridge, they can only be selected on the cartridge parameter pages.

Manage locations

For business flows to run correctly, each environment must have a number of locations associated with it, for example, source and target database locations and credentials, and a file location. Locations are initially specified during installation and automatically loaded into the Web Console. You can then view and manage all of these locations from the Web Console.

NOTE: If passwords are changed for any of your schemas or databases, you must edit the location and change the password.

TIP: While you can update location credentials within Web Console, to ensure security, you cannot change the encryption key for an environment from the Web Console. To change the encryption key, use the Password Manager utility. For more information, see [Change encryption keys and passwords, on page 185](#).

- [Default locations, below](#)
- [File archive settings, below](#)
- [Change locations and update passwords, on the next page](#)
- [Add locations, on page 86](#)
- [Parameterize locations, on page 87](#)
- [Configure Hadoop with HPE Structured Data Manager, on page 88](#)

Default locations

The following are the default locations in Web Console. You can modify and add to these locations.

| Location name | Description |
|-------------------|--|
| ARCHIVE_ACCESS_DB | Is the database location and credentials where the archive access schema resides. Only for database to database archiving. |
| DELETE_IDX_FS | Is the local file system for the delete index. |
| HISTORY_DB | Is the target (archive) database location and credentials. Only for database to database archiving. |
| HISTRELOC_DB | Is the database location and credentials where the relocation schema resides. Only for database to database archiving. |
| LOCAL_ARCHIVE_FS | Is the path where database to file stores archived XML or CSV files. |
| LOCAL_TEMP_FS | Is the path of the temporary file space. |
| OBTINTF_DB | Is the database location and credentials where the interface schema resides. |
| OLTP_DB | Is the source (active) database location and credentials. |

File archive settings

In **Business Flow Management > Settings > File Archive**, you will find the following file paths defined:

- **Archive Path** is the root path for your archive file storage. The path entered here can be referenced as `${archivePath}` elsewhere in the Web Console. For example, when defining your temporary directory, you might enter something like `${archivePath}/temp/` to indicate that you want to store temporary files in a subdirectory called `temp` inside of your root archive path. If you later decide to change the archive path, you do not need to update your temporary directory specification. It will update automatically because you used the variable for your archive path.
- **Default archive directory** is the location where database-to-file archive files are stored. The default value is `${archivePath}/archivedata/${environment}/`, which means that the path is the root archive path plus `archivedata` plus the current environment name.
- **Default temporary directory** is the location where Web Console stores temporary files. The default value is `${archivePath}/temp/${environment}/`, which means that the path is the root archive path plus `temp` plus the current environment name.
- **Default delete index directory** is the location where indexes are stored. The default value is `${archivePath}/delete_idx/${environment}/`, which means that the path is the root archive path plus `delete_idx` plus the current environment name.

Change locations and update passwords

If you need to change a location, such as the archive file location, or update a location's credentials, for example, the source database's password, you can do so from the Web Console.

To change a location or update passwords

1. Navigate to **Business Flow Management** in the Web Console.
2. Click **Environment**.
3. Click **Locations** from the left navigation pane.
4. Select a location.
5. Change the location properties as appropriate.

File system locations, such as `LOCAL_TEMP_FS` and `LOCAL_ARCHIVE_FS`, are typically expressed with variables:

- `${environment}` represents the name of the current environment.
- `${archivePath}` represents the archive path. You can change the archive path from **Business Flow Management > Settings > File Archive**. See [Change the archive path, on the next page](#).

5. Click **Apply**.
6. If the source database password for an environment changed, click **Environment > Admin Credentials** and update the password.
7. For database to database with a distributed instance, if you changed the interface schema or database password:
 - a. Log in to the active database as the interface user, for example, `OBT_IF`.
 - b. Drop the database link that points to the history relocation schema or database. This link was created when you deployed the business flow.
 - c. Re-create the database link with the same name.
 - d. Stop and start the Web Console.

8. For database to database with a distributed instance, if you changed the OLTP history relocation schema or database password:
 - a. Log in to the history database as the relocation user, for example, OBT_RELOC.
 - b. Drop the database link that points to the OLTP interface schema or database. This link was created when you deployed the business flow.
 - c. Re-create the database link with the same name.
 - d. Stop and start the Web Console.

NOTE: You can also change passwords using the password manager utility from the command line. For more information, see [Update database location passwords, on page 187](#).

Change the archive path

You can change the root archive path of Web Console and any locations based upon it (`${archivePath}`) will change, too.

To change the archive path

1. Navigate to **Business Flow Management** in the Web Console.
2. Click **Settings**.
3. Click **File Archive** from the left navigation panel.
4. Update the Archive Path. As soon as you make a change to this field, a radio button prompts you to choose between:
 - **Change for future archives** means you only want the path change to apply to future archiving operations. Existing archive files will remain in their current location.
 - **Move the data and apply for future archives** means you want to both move existing archive files to the new location and apply the path change to future archiving operations.

NOTE: If you choose **Move the data and apply for future archives**, any AQS collections that reference the old archive path will be automatically updated for you. See [Use AQS to access database to file output, on page 136](#) for more information about collections.

5. Change other settings as desired.
6. Click **Apply**.

Add locations

You may wish to create additional locations archiving. For example, in some cases you may want to archive files in a cloud location, such as HPE Cloud Object Storage, and, in other cases, you may wish to archive to a backend located on your network, such as Autonomy Consolidated Archive (ACA).

Database to file locations

Available location types for database to file are:

- ACA (On Premises)

TIP: For ACA, the timeout in the server's `ACAConnector.cfg` file under `[ViewServer]` should be increased from 60 seconds to 3600 (1 hour). 60 seconds is not enough time to extract large files.

- Amazon-S3
- AQS_Cache
- Centera
- Digital Safe
- Filesystem
- HDFS (Hadoop Filesystem), see
- HPE Cloud Object Storage
- SDM Gateway, see [Perform multiple data movements, on page 53](#)
- SSH
- XAM (deprecated)

NOTE: XAM only appears as an option if you have added `storage.xam.enable=true` to the `outerbay.properties` file, which is stored in the home directory under `config`. See [Configure the Web Console, on page 17](#).

- HPE Content Manager

Data access locations

Available location types for data access cartridges to query against are:

- AQS (Data Access Only)
- Vertica (Data Access Only)

Database to database locations

Available location types for database to database locations are:

- DB2
- JDBC URL
- Oracle
- PostgreSQL
- SQL Server
- Sybase

Parameterize locations

Within locations, you can make reference to business flow parameters, which allow you to dynamically set the path. For example, suppose that you want to include the runtime parameter values and the date in the path for your database to file extracts. You could enter the following for the **Directory** property of your file system location:

```
C:\myDirectory\${OA_CUST_ID}\${OA_ORD_YEAR}\csv\${DATE}
```

In this example, the parameters are indicated by `${}`. `OA_CUST_ID` and `OA_ORD_YEAR` are runtime parameters used by the business flow you are running. `OA_` is a prefix added to each parameter and the rest is the name given to the parameter when it was defined in Designer. `DATE` is a system parameter that inserts the current date at the time the business flow runs. This type of parameterization can be particularly useful when you want to separate your output files dynamically based on some set of parameters used during the business flow run.

TIP: You can reference any parameter that is referenceable from Groovy in the location definition as well. See the *HPE Structured Data Manager Developer's Guide* for more information about parameters you can reference from Groovy.

Configure Hadoop with HPE Structured Data Manager

HPE SDM accesses Hadoop as a Java client and requires two configuration files:

- `core-site.xml`
- `hdfs-site.xml`

In `outerbay.properties`, which is located in your application data directory, you must specify the path to these two configuration files, for example:

```
hadoop.config.dir=C:/SDM/HDFS
```

NOTE: You must use forward slash (/) as the directory separator even on Windows. If you use a backslash (\) in the `hadoop.config.dir` path, HPE Structured Data Manager throws an exception.

Refresh test environments

If you are maintaining a test environment in which you design and test your models, cartridges, and business flows, you may want to periodically refresh it from your production environment. For example, if the production schema has changed or been upgraded in some way, you need to reflect those changes in your test environment. You can duplicate changes to the production environment in your test environment without reinstalling HPE Structured Data Manager.

Prerequisites

For this process to work, your environments must meet the following criteria:

- Both environments must use the same version and configuration of HPE Structured Data Manager.
- Both environments must use the same version and configuration of Oracle. This procedure only applies to Oracle databases.

To refresh your test environment

1. In the test environment, clone the following databases from the production environment:
 - active (source) databases
 - repository databases
 - target databases (for the database to database distributed instance configuration only)

2. In the cloned databases, update any database links to use the test environment's host names, port numbers, and so on. Otherwise, the database links will continue to resolve to the production databases rather than the test databases.
3. Manually copy the product tree and the OBT_HOME directory. The OBT_HOME directory is defined in `<install_dir>/obt/config/obt.env`.
4. In the replicated location, edit the `<install_dir>/obt/config/obt.env` to reset the OBT_HOME for the replicated environment, and make a copy of `obt.env` as `env2.obt.env`.
5. Check that the `OBT_HOME/config/directories.properties` file has the correct settings for the replicated environment.
6. Navigate to the following directory:
`<install_directory>\obt\foundation\components\replicate`
 where `<install_directory>` is the location where you installed HPE Structured Data Manager.
7. Create a text file for the environment properties. The default name for this file is `replicate.properties`, which is defined in `replicate_environment.xml` as `<property name="properties.file" value="replicate.properties"/>`.

The file must contain the following properties:

| Property | Description |
|---|--|
| <code>env1.repository.password</code> | The repository password for the original environment. |
| <code>env1.encryption.key</code> | The encryption key for the original environment. |
| <code>env2.encryption.key</code> | The encryption key to be used for the duplicate environment. |
| <code>env2.product.path</code> | The path to be used fro the duplicate environment. |
| <code>env2.source.dbadmin.password</code> | The database administrator password for the source database. |
| <code>env2.source.dbadmin.userid</code> | The database administrator username for the source database. |
| <code>env2.hist.dbadmin.password</code> | The database administrator password for the target database. (distributed instance only) |
| <code>env2.hist.dbadmin.userid</code> | The database administrator username for the target database. |
| <code>env2.repository.password</code> | The repository password for the duplicate environment. |
| <code>env2.interface.password</code> | The password for the interface user. |
| <code>env2.transparency.layer.password</code> | The password for the OBT_BUILD_AA user. |
| <code>env2.transparency.layer.location</code> | The location of the OBT_BUILD_AA user. Valid values are source or target. |

| Property | Description |
|------------------------------------|---|
| env2.relocation.interface.password | The password for the OBT_RELOC user. |
| env2.source.password | The password for the interface user (for example, obt_if) on the source database. |
| env2.source.host | The host name for the source database. |
| env2.source.server | The source database name. |
| env2.source.port | The port number for the source database (for example, 1522). |
| env2.repo.dbadmin.userid | The database administrator username for the repository database. |
| env2.repo.dbadmin.password | The database administrator password for the repository database. |
| env2.repo.host | The host name for the repository database. |
| env2.repo.server | The repository database name. |
| env2.repo.port | The port number for the repository database (for example, 1522). |
| env2.environment | The new environment name of the test if it is different from the production environment. |
| env2.hist.host | The host name for the archive (target) database. |
| env2.hist.port | The port number for the archive database (for example, 1522). |
| env2.hist.server | The archive database name. |
| env2.hist.password | The password for the relocation user (for example, obt_reloc) on the archive database. |
| env2.hist.schema.password | The password for the history schema. |
| env2.aa.schema.password | The password for the archive access schema. |
| replace.dblink | Whether to replace database links (true or false). |
| env2.oltp.to.hist | If you set <code>replace.dblink</code> to true, the new name of the database link from the active database to the archive database, for example, OBT_OLTP_TO_HIST2. |
| env2.hist.to.oltp | If you set <code>replace.dblink</code> to true, the new name of the database link from the archive database to the active |

| Property | Description |
|-------------------|---|
| | database, for example, OBT_HIST_TO_OLTP2. |
| env2.aa.to.remote | If you set <code>replace.dblink</code> to true, the new name of the database link from the archive access schema, for example, OBT_AA_TO_REMOTEDB2. |
| environment.id | The name of the old environment from which this procedure is refreshed. |
| configuration | The environment's configuration: <ul style="list-style-type: none"> • OA means the environment has only database to file. • LA means the environment has database to file and database to database. |
| configuration.la | If the environment includes database to database, then this property indicates: <ul style="list-style-type: none"> • SI for single instance, where source and target are in the same database • DI for distributed instance, where target is in a separate database) |
| recreate.views | If you set to true, it recreates the views in the archive access schemas, replacing the database link name with the new name in the replicated database environment. This is only relevant for a distributed instance configurations, and might be very time consuming. If set to false, it will not recreate these views. |
| replace.db_link | <p>If you set to true, insert a new database link name in the following properties to drop and recreate the database links with the changed database domain:</p> <ul style="list-style-type: none"> • env2.oltp.to.hist • env2.hist.to.oltp • env2.aa.to.remote <p>NOTE: The DBLINK name should be the same as the global name of the target if the Oracle GLOBAL_NAMES property is set to true. Furthermore, suffixes on the global name will be applied to DBLINK, which can cause issues on the target database if it has GLOBAL_NAMES set to true and</p> |

| Property | Description |
|---------------------|---|
| | a different suffix. |
| repository.username | The repository user name, for example, obt_rep. |

Database to file environment example

```
env1.repository.password=obt_rep2
env1.encryption.key=q2

env2.encryption.key=q2
env2.product.path=E:\\program files\\HPESDM740\\obt
env2.source.dbadmin.password=manager
env2.source.dbadmin.userid=system
env2.repository.password=obt_rep2
env2.interface.password=obt_if2
env2.source.password=obt_if2
env2.source.host=localhost
env2.source.server=orcl2
env2.source.port=1522
env2.repo.host=localhost
env2.repo.server=orcl2
env2.repo.port=1522

environment.id=OA
configuration=OA
configuration.la=DI
repository.username=obt_rep
```

Database to database environment example

```
env1.repository.password=obt_rep2
env1.encryption.key=q2

env2.encryption.key=q2
env2.product.path=E:\\program files\\HPSDM720\\obt
env2.source.dbadmin.password=manager
env2.source.dbadmin.userid=system
env2.hist.dbadmin.password=manager
env2.hist.dbadmin.userid=system
env2.repository.password=obt_rep2
env2.interface.password=obt_if2
env2.transparency.layer.password=obt_build_aa2
env2.transparency.layer.location=source
env2.relocation.interface.password=obt_reloc2
env2.source.password=obt_if2
env2.source.host=localhost
env2.source.server=orcl2
env2.source.port=1522
env2.repo.host=localhost
env2.repo.server=orcl2
env2.repo.port=1522
env2.hist.host=localhost
env2.hist.port=1522
env2.hist.server=orcl3
env2.hist.password=obt_reloc2
env2.hist.schema.password=target2
env2.aa.schema.password=aa2

replace.dblink=true
env2.oltp.to.hist=OBT_OLTP_TO_HIST2
env2.hist.to.oltp=OBT_HIST_TO_OLTP2
env2.aa.to.remote=OBT_AA_TO_REMOTEDB2
environment.id=LADI
configuration=LA
configuration.la=DI
repository.username=obt_rep
```

8. Run the appropriate script for your operating system:

| Platform | Syntax |
|----------|---|
| UNIX | <install_dir>/obt/bin/replicate_environment.sh |
| DOS | <install_dir>\obt\bin\replicate_environment.bat |

NOTE: Remotely copying the production environment to another server on UNIX is not supported.

9. To replicate more than one environment from the same production instance, repeat [Manually copy the product tree and the OBT_HOME directory](#). The OBT_HOME directory is defined in <install_

[dir>/obt/config/obt.env.](#), on [page 89](#) through [Run the appropriate script for your operating system](#), on the [previous page](#) with the following modifications:

- The repository password and encryption key of the first environment should be changed to that of the new environment because they will have already been updated in the first run of `replicate_environment.bat`.
 - If the database links have to be renamed, set `replace.dblink` to `true` and provide the new link names.
10. Update the `sqltuning.properties` file, which can be found changing references to the old environment name to the new environment name.
 11. Restart the Web Console.

Rename environments

After you refresh the test environment, you may want to change the environment name.

To change the name of an existing environment

```
rename_environment.bat/sh -e [old_env_name]  
-n [new_env_name]
```

To change the environment description

```
rename_environment.bat/sh -e [old_env_name]  
-n [new_env_name] -d [new_env_description]
```

Chapter 10: Manage business flows

Before you run business flows, you should perform some checks; for example, confirming parameter settings. Beyond that, certain types of business flows require that you perform preparatory tasks before running them.

After you run your business flows, you need to access the archived data or possibly recover from error conditions.

TIP: Step-by-step information for the Web Console user interface can be found in the Web Console online help system.

- [Deploy business flows](#)
- [Prepare to run business flows](#)
- [Recover or cancel business flows and jobs](#)
- [Access archived data](#)

Deploy business flows

Business flows contain the activities that actually perform the movement, copying, and reloading of data, as well as other business logic (Groovy scripts). In order to run a business flow, you must first deploy and configure it in an environment. You can deploy business flows either from Designer or from the Web Console.

To deploy business flows using the Web Console, they must first be generated by you or someone else in Designer, and the resulting deployment files placed in a location accessible to the Web Console.

NOTE: Business flows containing 2-tier configuration must be deployed separately from business flows containing 3-tier configuration.

NOTE:
Deployment of Database to Database business flow is a prerequisite for the deployment of a Database to File business flow in the 3-tier mode.

Generated business flows are zip files with a .busflow extension. By default, business flows are stored in one of the following locations, but you can import .busflow files from other locations as well:

| Location | Description |
|---|---|
| businessflow in home directory | Business flows stored in the main business flow directory can be used by all users. |
| businessflow/<env_name> in home directory | Business flows stored in an environment directory can only be accessed by users who are assigned to that environment. |

For more information about the home directory, see [Configure the Web Console, on page 17](#).

The deployed business flows are stored in the following directory under the home directory:

/artifacts/<environment_name>/businessflow/<busflow_name>/

| Parameter | Description |
|------------------|--|
| environment_name | The name of the environment the business flow was deployed into. |
| busflow_name | The name of the business flow. |

If you need to modify your business flow and its components to account for changes to the database, or for business or legal requirements, then you must modify the business flow or its components in Designer and redeploy it.

As with a new deployment, if you plan to redeploy business flows using the Web Console, they must first be generated by you or someone else in Designer, and the resulting deployment files placed in a location accessible to the Web Console.

The redeployed business flows are stored in the same location as when they were originally deployed.

When you no longer need a particular business flow or data access cartridge, you can use the Web Console to remove (uninstall) it.

CAUTION:

You should never remove the deployment directory or its contents manually from the file system. It must be properly uninstalled from the Web Console. Otherwise, you will get unexpected results when you try to redeploy, run, or uninstall the business flow.

See the Web Console online help and [Deploy and run business flows, on page 43](#) for more information about deploying business flows in the Web Console.

Prepare to run business flows

Some business flows, such as those that reload data, require some preparation before running.

Prepare to reload archived data

Before running a reload from database business flow, you need to ensure that the same indexes were created on the archive database that existed on the active database. Otherwise, the performance of the reload will be adversely affected.

If you ran your archive jobs with the default parameter settings for creating indexes on history tables, then the indexes were created at that time and you can proceed with your reload business flow without taking any further action. If the parameter was set otherwise and indexes were not created during the archive business flow, you can create the indexes manually yourself. You can also rerun the archive business flow with a parameter value that causes HPE Structured Data Manager to create the indexes on the archive database.

TIP: If the archiving business flow was previously run with the Type of indexes to create on History parameter set to Operate on all indexes, the indexes should have been created at that time and you need not perform this procedure.

To create the indexes on the history tables using Web Console

1. Navigate to **Business Flow Management** in the Web Console.
2. Navigate to the Parameters - Database to Database page in the Web Console, **TASKS > System Parameters > Database to Database**.
3. Set the value of the **Type of indexes to create on History** parameter to **Operate on all indexes**.
4. Rerun the business flow using parameter settings that will not archive any data. For example, you could select a data range that does not match your retention policies.

The indexes are automatically created on the archive database, and you can now proceed to reload your archived data.

Prepare to upload archived data

Before you upload your archived data, you should consider the following:

- See [Map schemas by cartridge, on page 113](#) if you want to map the archived data into a schema name that is different from the source.
- The upload user must have create table, drop table, create index, and drop index privileges on the target database where the data is being uploaded.
- If you are planning to upload from an archive file generated against Oracle to SQL Server, you should explicitly set the precision and scale on numeric columns, where appropriate, in the source database. Otherwise, if SQL Server has more restrictive precision and scale settings than Oracle, the upload may fail.
- The default batch size for upload is 1000 rows. HPE Structured Data Manager will store that many rows per worker prior to inserting them into the database. For example, if the number of workers is set to 5, then 5 threads will be running with each holding 1000 rows in memory before inserting into the target database. If you need to adjust the batch size, change `upload.batch.size` in `outerbay.properties` in `config` under the home directory.

Recover or cancel business flows and jobs

In general, it is best practice to never leave a business flow or job in one of the following states:

- Failed
- Suspended
- Not Responding

Otherwise, you may get unexpected results when you attempt to run the same or other, related business flows, or re-deploy the business flow.

TIP: To check the status of a business flow or job, go to **Business Flow Management > TASKS > Monitoring > Overview** in the Web Console.

If a business flow does not complete successfully, you can use the Web Console to:

- **Recover it.** Recovering the job or business flow reruns it using the same parameters as the first run. If possible, recovering is the best option.
- **Cancel it.** Cancelling the job or business flow should be your option of last resort. If you cannot

recover for some reason and you need to redeploy the business flow, you can cancel the business flow. Cancelling updates the status of the business flow or job, but does not clean up any of the selection tables nor does it remove partially copied or deleted data. You may end up with unnecessary information in your database, which can never be used after canceling the job, and it will continue to occupy space. Consequently, cancelling a job can lead to inconsistent results.

NOTE: If you cancel an upload business flow, it can result in indexes not being created on the target database, which can lead to performance problems.

IMPORTANT:

If a business flow containing an advanced selection database to file cartridge does not complete successfully, see the *HPE Structured Data Manager Troubleshooting Guide* for further information before rerunning the business flow.

Access archived data

After you archive your inactive data to an archive database or file, you still need to access it. From the Web Console, you can perform a variety of operations to facilitate access to the archived data:

- For database to database archives, you may enable transparency to allow users to see the archived data as though it were still in the active database.
- For database to file archives, you may relocate, copy, or delete the archived data from the Web Console.

This section describes:

- [Create database to database archive access transparency, below](#)
- [Relocate database to file archive data, on the next page](#)
- [Access the archive data file, on page 100](#)

Create database to database archive access transparency

Data transparency refers to the ability of users to see archived data as though it were still in the active database. If you implement data transparency, both active and archive data can be accessed through the current user interface. All data relationships necessary to maintain data and application integrity are retained.

After archiving your data, you can run the Create Archive Access job to enable transparency, which allows your users to view your archived data directly from the application.

NOTE: Contact your HPE sales representative for information about licensing requirements for data transparency.

The Create Archive Access job creates or recreates objects within the Archive Access schemas or databases. The running time of the script depends on the number of cartridges included in your business flows, the number of managed tables, and the number of views, stored procedures, and packages that are dependent on those objects. To customize this job, see the *HPE Structured Data Manager API Reference Guide*.

You should run the Create Archive Access job whenever any of the following occur:

- Changes are made to tables, views, stored procedures, or packages in the active database.
- New cartridges are installed.
- Cartridges are upgraded.
- New patches are installed on the active database.
- Any of the ArchiveAccessConfiguration groovy scripts are run.

See the *HPE Structured Data Manager API Reference Guide*

NOTE: The Archive Access schema is unstable while the Create Archive Access job is running. To avoid errors, wait until after the job has completed successfully before accessing the archive.

1. Navigate to **Business Flow Management** in the Web Console.
2. Confirm that the configuration mode parameter is properly set by navigating to **TASKS > Monitoring > System Parameters > Database to Database**.
See [Manage environments, on page 78](#)
3. Click **TASKS > Launch > Business Flow**.
4. Click **Create Archive Access**. The job launch page for the Create Archive Access job displays.
5. Click **Run** to run the job immediately, or **Schedule** to schedule the job for a later time.
6. Click **Confirm** to confirm the job. When the job completes successfully, you should be able to access the archived data from your application.
7. Test accessing some of the archived data from your application to ensure that the archive access is working.

NOTE: If you have uploaded the same data to the same database, you can have duplicate data because the upload process only checks for duplicate data when uploading lookup and MTU data. (It checks the target database for duplicate data.) To prevent duplicate data, the reload process checks for duplicate data based on the data movement key of the table. If the table does not have a data movement key or the data movement key is ROWID, reload may result in duplicate rows.

Relocate database to file archive data

After you archive from database to file, you use **TASKS > Launch > Relocate** to relocate or copy the archived data to a variety of locations:

- source database from which it was archived
- another database
- another file location

You can also delete the archived data from a backend.

For more information about relocating files, see the online help.

Reload and upload

When uploading or reloading the archived data, the rows are inserted into the database in batches. The size of these batches of inserts is set by the `upload.batch.size` property in `outerbay.properties`,

which is located in `config` under the home directory. See [Configure the Web Console, on page 17](#).

A database commit can be performed after each batch of inserts from the file. Performing the database commit at the batch level ensures that the size of the rollback segment remains smaller, but referential integrity cannot be guaranteed until reload completes and all of the data is committed.

For all cases of upload (XML and CSV), only batch commits are performed.

For reload from an XML file, only file group commits are performed. That is, after inserting from an entire group of files, a database commit is performed.

For reload from a CSV file, you can control when the database commit is done. Use the following setting in `<install_dir>\obt\config\outerbay.properties` to control database commits for reloading from a CSV file:

- `reload.batch.commit=true` indicates batch level database commits. This setting has the benefit of smaller rollback segments, but it can lead to referential integrity issues until reload completes and all of the data is committed for every table.
- `reload.batch.commit=false` is the default setting and indicates file group level database commits. It can lead to very large rollback segments and, if the reload should fail, it will restart for the entire file group.

Access the archive data file

To access contents of your archive data files, you use the Archive Query Server. See [Use AQS to access database to file output, on page 136](#).

Chapter 11: Run business flows and jobs from the command line

To provide additional flexibility, HPE Structured Data Manager enables you to run business flows and jobs from the command line as well as from the Web Console. This capability is useful for environments where you cannot easily access the Web Console or need to script the running of business flows or jobs.

- [Before you begin](#)
- [Run business flows](#)
- [Recover and cancel business flows and jobs](#)
- [Access the archived data](#)

Before you begin

Before you run business flows, you should perform some checks; for example, confirming parameter settings. Beyond that, certain types of business flows require that you perform preparatory tasks before running them. In particular, business flows that involve reloading data or bulk archiving require some preparation before running:

- The preparatory tasks for running business flows from the Web Console also apply to running on the command line. See [Manage business flows](#).
- To run business flows or jobs from the command line, you may need some or all of the following information, which you can obtain from reviewing the deployed business flows in the Web Console:
 - Environment name
 - Environment ID
 - Business flow name
 - Job parameters
 - Job name
 - Job run ID
 - Destination location

Run business flows

To run business flows or jobs from the command line

1. Make sure that you have performed the relevant tasks from [Before you begin, above](#).
2. Open a command window.
3. Navigate to the following directory:

`<install_directory>/obt/bin`

where `<install_directory>` is the location where you installed the software.

4. Run the business flow using the appropriate syntax:

| Platform | Syntax |
|----------|--|
| UNIX | <code>./launch_businessflow.sh [-r] -e <environment_name> -j <business_flow_name> <parameters></code> |
| DOS | <code>launch_businessflow.bat [-r] -e <environment_name> -j <business_flow_name> "<parameters>"</code> |

| Parameter | Description |
|--------------------|--|
| -r | The recover flag, -r, allows you to recover jobs and business flows. See Recover business flows, on the next page . Although the recover flag, -r, is optional, HPE recommends that you use it whenever you run a job or business flow. |
| environment_name | The name of the environment. This is optional if you only have one environment. |
| business_flow_name | The name of the business flow. |
| parameters | The parameters as a space-separated list with each entry in the form <name>=<value>. For Windows, each individual entry must be enclosed by quotation marks. The parameter values for all cartridges included in the business flow are required. For a date or time policy parameter, use the following format: [YY]YY.MM.DD[HH24:MI:[SS[.SSS]]] |

TIP: To verify the usage of the job from the command line, enter `launch_businessflow.bat -h`.
The usage is displayed.
To verify the parameters for the job, enter `launch_businessflow.bat -j <business_flow_name> -e <environment_name> -h`.
After you enter the encryption key, the parameters are displayed.

DOS example

```
launch_businessflow.bat -e Oracle_env -j orders_flow_d2d -r "Min_Months_to_Retain=24"
```

TIP: If the parameter value is left blank, and a default value was defined in Designer, the default value is used. If there is no default value, NULL is used.

NOTE: Running an undo business flow requires the Run ID for the business flow you want to undo. The Run ID for the business flow is displayed on the Run Summary page.

5. Enter the encryption key at the prompt. The encryption key is case sensitive.
The job runs, and displays a success or failure message.

6. Optionally, launch the Web Console to monitor the progress of the job.

Recover and cancel business flows and jobs

- [Recover business flows](#)
- [Recover jobs using the recovery script](#)
- [Cancel jobs from the command line](#)

Recover business flows

Using the recover flag, `-r`, allows you to rerun a business flow using the same parameters as the first time it was run. You can also use the `-r` flag to continue the business flow after a scheduled pause activity.

NOTE: To recover a job, the job parameters must be exactly the same as the last time you ran it. If there are any differences in the syntax, even an extra space, a new job is launched instead of the previous job being recovered.

1. Make sure that you have performed the relevant tasks from [Before you begin, on page 101](#).
2. Open a command window.
3. Navigate to the following directory:
`<install_directory>/obt/bin`
where `<install_directory>` is the location where you installed the software.
4. Enter the same job name and parameters as the job you want to recover.

DOS example

If you ran the business flow example given in [Run business flows, on page 101](#), you would run:

```
launch_businessflow.bat -e Oracle_env -j orders_flow_d2d -r "Min_Months_to_
Retain=24"
```

5. Optionally, launch the Web Console to monitor the progress of the job.

Recover jobs using the recovery script

Using the recovery script to recover jobs enables you to select the specific job you want to recover.

To recover a job using the recovery script

1. Make sure that you have performed the relevant tasks from [Before you begin, on page 101](#).
2. Locate the Group ID of the failed job you want to recover.
The Group ID is displayed on the Web Console History page.
3. Open a command window.
4. Navigate to the following directory:
`<install_directory>/obt/bin`
where `<install_directory>` is the location where you installed the software.

5. Run the appropriate recovery script for the job you want to recover:

| Job | Platform | Syntax |
|---|----------|---|
| Any business flow or database to file job | UNIX | <code>./recover_job.sh -e <environment_name> -g <group_ID></code> |
| | DOS | <code>recover_job.bat -e <environment_name> -g <group_ID></code> |
| The Create Archive Access job | UNIX | <code>./recover_create_aa.sh -e <environment_name> -g <group_ID></code> |
| | DOS | <code>recover_create_aa.bat -e <environment_name> -g <group_ID></code> |

| Parameter | Description |
|------------------|---|
| environment_name | The name of the environment. This is optional if you only have one environment. |
| group_ID | The Group ID of the failed job as it appears in the Web Console. |

DOS example

```
recover_job.bat -e "Oracle_env" -g "3"
```

6. Enter the encryption key at the prompt. The encryption key is case sensitive.
The job runs, and displays a success or failure message.
7. Optionally, launch the Web Console to monitor the progress of the job.

Cancel jobs from the command line

Cancelling jobs from the command line requires the `cancelJob` API Groovy script. The script allows you to update the status of a job to Cancelled. However, you should only cancel jobs in extreme cases because using the `cancelJob` API does not restore the state of the database to beginning of a failed job. You will have to manually restore the state of the database.

To cancel a job

1. Locate the Group ID of the business flow you want to cancel.
The Group ID is displayed on the Web Console History page.
2. Navigate to the following directory:
`<install_directory>/obt/bin`
where `<install_directory>` is the location where you installed the software.

3. Run the `cancelJob` API using the appropriate syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <code>./launch_groovyscript.sh -e <environment_name> -f ../scripts/cancelJob.groovy <groupRunId></code> |
| DOS | <code>launch_groovyscript.bat -e <environment_name> -f ..\scripts\cancelJob.groovy <groupRunId></code> |

where

`<groupRunId>` is the Group ID as defined on the Web Console.

`<environment_name>` is the name of the environment.

4. Optionally, start the Web Console to monitor the progress of the job.

Access the archived data

- [Create database to database archive transparency](#)
- [Relocate database to file archive data](#)

Create database to database archive transparency

See [Create database to database archive access transparency](#)

To create data transparency from the command line:

1. Ensure that you have performed the relevant tasks from [Before you begin, on page 101](#).
2. If necessary, perform the following tasks to ensure your configuration mode is set correctly for the transparency layer:
 - a. Navigate to **Business Flow Management** in the Web Console.
 - b. Navigate to the Parameters - Database to Database page.
 - c. Check the Archive Access section to ensure that the configuration mode parameter is set to the configuration option you want to use.

See [Manage environments, on page 78](#).

3. Run the Create Archive Access job using the following syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <code>./launch_create_aa.sh [-r] -e <environment_name> [<transparency_layer_name>]</code> |
| DOS | <code>launch_create_aa.bat [-r] -e <environment_name> [<transparency_layer_name>]</code> |

| Parameter | Description |
|------------------|--|
| -r | The recover flag, -r, allows you to recover jobs and business flows. See Recover business flows . Although the recover flag, -r, is optional, HPE recommends that you use it whenever you run a job or business flow. |
| environment_name | The name of the environment. This is optional if you only have one environment. |

DOS example

```
launch_create_aa.bat -e Oracle_Environment -r
```

4. Enter the encryption key at the prompt. The encryption key is case sensitive.
The job runs, and displays a success or failure message.
5. Optionally, launch the Web Console to monitor the progress of the job.
When the job completes successfully, you should be able to access the archived data from your application.
6. Try accessing some of the archived data from your application to ensure that the archive access is working.

Relocate database to file archive data

See [Relocate database to file archive data, on page 99](#)

To relocate database to file archive data from the command line:

1. Make sure that you have performed the relevant tasks from [Before you begin, on page 101](#).
2. Open a command window.
3. Navigate to the following directory:
`<install_directory>/obt/bin`
where <install_directory> is the location where you installed the software.

4. Run the job from the command line using the following syntax:

| Platform | Syntax | |
|----------|---|--|
| UNIX | <code>./launch_ea_job.sh [-r] -e <environment_name> -j <jobName></code> <code>[sourceJobRunID=<runID>] [sourceLocation=<source_loc>]</code> <code>[summaryFileName=<file1,file2,...>] [destLocation=<dest_loc>]</code> <code>[mappingFileName=<map_file>]</code> | |
| DOS | <code>launch_ea_job.bat [-r] -e <environment_name> -j <jobName></code> <code>[sourceJobRunID=<runID>] [sourceLocation=<source_loc>]</code> <code>[summaryFileName=<file1,file2,...>] [destLocation=<dest_loc>]</code> <code>[mappingFileName=<map_file>]</code> | |

| Parameter | Description | |
|------------------|---|---|
| -r | The recover flag, -r, allows you to recover jobs and business flows. | |
| environment_name | The name of the environment. This is optional if you only have one environment. | |
| jobName | One of the following: | |
| | XML_ARCHIVE@COPY_BE_TO_DB | Reloads data from a backend to the source database. |
| | XML_ARCHIVE@COPY_BE_TO_BE | Copies data from a backend to a backend. |
| | XML_ARCHIVE@UPLOAD_BE_TO_DB | Uploads data from a backend to a different database than the original source, checks for missing schema objects, and creates any schema objects that are missing. |
| | XML_ARCHIVE@MOVE_BE_TO_DB | Reloads data from a backend to the source database, and then deletes it from the backend. |
| | XML_ARCHIVE@MOVE_BE_TO_BE | Copies data from one backend to another backend, and deletes it from the original. |
| | XML_ARCHIVE@DELETE_FROM_BE | Deletes the data from the backend. |
| | XML_ARCHIVE@DELETE_FROM_DB | Deletes the data from the database. |
| | XML_ARCHIVE@REGISTER_ARCHIVE | Registers a file archive with another HPE Structured Data Manager instance and creates new lineage with the destination as a file system in the current HPE |

| Parameter | Description | |
|-----------------|---|-----------------------------------|
| | | Structured Data Manager instance. |
| runID | The run ID of the job that placed the archive at the current location. | |
| source_loc | The Web Console location where the file archive to import is stored. | |
| file1,file2,... | A list of summary files for the file archives you wish to import. | |
| dest_loc | The name of the destination location as defined in the Web Console. Required for COPY_BE_TO_BE, MOVE_BE_TO_BE, and UPLOAD_BE_TO_DB only. | |
| map_file | The name of the mapping file to use during upload when disconnected from the source database. See Map schemas by cartridge, on page 113 for more information. | |

DOS examples

Copy a file archive from one backend to another:

```
launch_ea_job.bat -e "MyEnvironment" -j "XML_ARCHIVE@COPY_BE_TO_BE" -r
"sourceJobRunId=3" "destLocation=LOCAL_ARCHIVE_FS"
```

Import a file archive leaving the original intact:

```
launch_ea_job.bat -e oracle_env -j XML_ARCHIVE@IMPORT -r "sourceLocation=Copy_
BE_2_BE" "summaryFileName=Order_Table_0955wwx5r66em90yntcmf46jj3ucudsscua_
summary.xml.gz" "destLocation=Import_From_FS"
```

Import a file archive deleting the original:

```
launch_ea_job.bat -e oracle_env -j XML_ARCHIVE@IMPORT_MOVE -r
"sourceLocation=Copy_BE_To_BE" "summaryFileName=Order_Table_
0hrqqo3q6cve036sebzj2kh1v3o8fr9o0zm_summary.xml.gz,Order_Table_
0hrqqo3q6cve09bfkjqx11ym6rm8phu7z7m_summary.xml.gz" "destLocation=Import_To_FS"
```

5. Enter the encryption key at the prompt. The encryption key is case sensitive.

The job runs, and displays a success or failure message.

6. Optionally, start the Web Console to monitor the progress of the job.

Chapter 12: Discovery

Personal data (personally identifiable information (PII) or sensitive personal information (SPI)) is information that can be used on its own or with other information to identify, contact or locate a single person, or to identify an individual in context. Any information that can be used to distinguish one person from another or could be used for de-anonymizing anonymous data is also considered personal data.

HPE SDM Discovery feature

The Discovery feature in HPE SDM enables you to define discovery projects that analyze, declare and manage sensitive data. The discoveries based on business needs can be grouped together into projects.

The Discovery Extension uses the database connection (OLTP_DB location) from the HPE SDM environments to retrieve the list of source tables and to scan the content of the tables.

NOTE:

The Discovery project is supported only for PostgreSQL repositories and requires the **obtdiscovery** schema in the SDM repository. The schema is created automatically upon running the extension for the first time.

Create a discovery project

To create a discovery project

1. Click **New Project**.
The New Project dialog opens.
2. Enter a project name in the **Project Name** box.

NOTE:

A project name is required. The project name must be unique.

3. Enter an optional description for the project in the **Description** box.
4. Select an environment from the **Environment** list.
5. Click **Save**.

NOTE:

To import a project created on a different SDM instance, see the section on [Exporting and Importing Projects](#).

The Project dashboard page for the new project opens, displaying the following tabs:

- **Overview.** Displays a summary of the project, including a donut graph of the records by classification.

| Button/Field | Description |
|-----------------------------|--|
| Project Options list | Select project options from the list: <ul style="list-style-type: none">◦ Edit Project. Select to edit an existing project. |

| Button/Field | Description |
|--------------|--|
| | <ul style="list-style-type: none"> ◦ Scheduled Scans .Click Scheduled Scans to view the scans scheduled to run for the discovery project. ◦ Export.. A locked project can be exported to allow transport to another SDM instance. ◦ Delete Project. Select to delete a project. |

- **Tables**. Displays a summary of the tables that are added to the project and provides the facility to register and scan the additional tables.

| Button/Field | Description |
|-------------------------|---|
| Home Button | Returns the list of projects. |
| Registration | <p>Add Tables – Click Add Tables to add additional tables to the discovery project from the registered schemas.</p> <p>Registered Schemas – Once you have added tables to the discovery project, click Registered Schemas to view the list of schemas registered to the discovery project.</p> <p>Add Schemas – Click Add Schemas to add schemas to the discovery project.</p> |
| Queue Processing | <p>Scan Now – Click to scan the table data.</p> <p>Schedule – Click to schedule a scan of the table data.</p> <p>Reset All – Displayed only when a table scan has been completed. Click Reset ALL to queue the table again for processing.</p> |
| Fields | Click to see all of the fields in the project and to review them as a whole. |
| Refresh | Click to refresh the page, or to select a refresh page interval. |
| Actions | <p>The following Actions applies to the selected table:</p> <ul style="list-style-type: none"> ◦ Scan Now – Click to scan the table data. ◦ Schedule Scan – Click to schedule a scan of the table data to happen at a later point. ◦ Reset – Click Reset to queue the table again for processing. ◦ Filter – Click Filter to add the SQL WHERE clause to the table data. ◦ Duplicate – Click Duplicate to duplicate the table data in the discovery project . |

- **Table Detail Page**. The Table Detail tab displays a detailed report on the tables of the project.

| Button/Field | Description |
|-------------------|--|
| Up a Level | Returns to the list of Tables . |
| Review | After the tables are scanned and the data are ready for review, select the columns of interest and and click Review to perform one of the following |

| Button/Field | Description |
|----------------------|--|
| | review actions: <ul style="list-style-type: none"> ◦ Confirm - To confirm the status of the column. ◦ Comment - To add a comment to the column. ◦ Override - To change the personal data type. ◦ Declare a false positive - To change the status of the column to non-sensitive. |
| Table History | Displays the scan history of the current table. |

- **Column Detail Page.** The Column Detail tab displays a detailed report on the columns of the project.

| Button/Field | Description |
|-------------------|---|
| Up a Level | Returns to the Table Detail page. |
| Review | After the tables are scanned and the data are ready for review, select the columns of interest and click Review to perform one of the following review actions: <ul style="list-style-type: none"> ◦ Confirm - To confirm the status of the column. ◦ Comment - To add a comment to the column. ◦ Override - To change the personal data type. ◦ Declare a false positive - Changes the status of the column to non-sensitive. |
| View Data | If the system is configured to allow data viewing and if you are granted access to the View Data role, then you can view the samples of data fetched from the source database to assist reviewing the column. |

- **Report.** The Report tab displays a detailed report on the current status of the project.

| Button/Field | Description |
|--------------------|--|
| HTML button | Click to export the report in HTML format. |
| PDF button | Click to export the report in PDF format. |

Exporting and Importing Projects

The Discovery module allows the export of projects so that they can be transferred to another instance of SDM with the discovery module enabled. You must also ensure that the version of the discovery module is same between the origin and destination.

NOTE:

The SDM repository user requires the `SUPERUSER` role in Postgres to enable the export or import

option.

Export Projects

1. Navigate to the **Project Overview** tab of the project you wish to export.
2. Expand **Project Options**, and select **Export**.
Export only becomes available after you have locked your project indicated by a lock icon next to the project name (see Edit Project for details on locking projects).
3. Click the **Export** button on the popup dialog to start the download, and your project will be downloaded in a zip file. You must ensure not to edit or rename this file.

Import Projects

On the main dashboard or projects tab, do the following:

1. Click the **New Project** button.
2. Click the **Import Project** switch.
3. Locate your exported project file using the **Choose file** button to open a browse dialog, click **Open** and then **Save**.

During the import, if an environment exists that matches the project environment, then the project is attached to this environment. If the environment does not exist, then the project is imported as remote.

Remote Projects

Remote projects are the projects that are collected by different instances of SDM, they are exported from that instance and imported to another instance. This allows reporting to be done remotely from the scanning.

Some features are disabled for remote projects such as adding tables, performing or scheduling scans. But you can still review the collected meta data and the review columns.

Chapter 13: Map schemas for upload or AQS cache

When uploading from a file to a database, you may have cause to perform some level of mapping. You may need to upload into a different schema name than the one from which you archived the data, which would require a simple mapping of one schema name to another. For example, if you archived from a SQL Server database and are now uploading to an Oracle database, you may prefer a different schema name.

In other cases, where you are uploading to or from non-native databases, you may need to perform a more elaborate mapping that includes data types. For example, you may have a `real` data type in your source that should be mapped to the `float` data type in an upload to SQL Server and the `NUMBER` data type in an upload to Oracle.

When using AQS cache, which uploads to PostgreSQL, you might also need schema mapping to account for differences between your source database and PostgreSQL. For example, your source database might allow longer name lengths than PostgreSQL.

NOTE: You also should be conscious of different length definitions. One database might measure length in characters while another uses bytes. For example, if you are uploading data that originated in a database that measured length in characters to one that measures length in bytes, you need schema mapping to reconcile those differences during upload.

NOTE:

In case of Oracle, if you have 'CHAR' column types, you must check the addendum "Extracting binary zero in Oracle CHAR or VARCHAR" to avoid issues sending data to AQS having char(0) ASCII data in char type columns of any table.

HPE Structured Data Manager provides two methods for performing such schema mappings.

- [Map schemas by cartridge](#)
- [Map schemas by database](#)

Map schemas by cartridge

One way to map schemas for upload is to create a mapping for your cartridge. This method is best for one-off mappings that you only need for a particular cartridge. If you are doing many, similar uploads, it may be more cost effective to create a database schema mapping. For more information, see [Map schemas by database, on page 116](#)

TIP: A schema mapping file can also be useful when uploading to a database where the maximum size of a column's data type is set smaller than the incoming data. The mapping file allows you to prevent errors or properly handle truncation, depending on the database you are using.

To create and use a schema mapping for your database to file cartridges

1. Create a schema mapping file in a location that is accessible to your HPE Structured Data Manager instance. The contents of the mapping file should adhere to one of the following formats, depending on whether you want to map:

- target catalog, schema, and table
- target catalog and schema
- target catalog only

The mapping of the qualified target occurs by matching the above conditions in the order they are specified in the mapping file. If there is a match to a condition, no further mappings are checked.

To map the target catalog, schema, and table, use the following format:

```
<src_catalog>.<src_schema>.<src_table>=<target_catalog>.  
<target_schema>.<target_table>
```

NOTE: In the case of Oracle, there is no catalog. Hence, you would omit the catalog in cases involving Oracle. For example, if you are going from Oracle to SQL Server, you would not need <src_catalog>.

NOTE: In case of Oracle, if you have 'CHAR' column types, you must check the addendum XXX to avoid issues sending data to AQS having char(0) ASCII data in char type columns of any table.

To map the target catalog and schema, use the following format:

```
<src_catalog>.<src_schema>=<target_catalog>.  
<target_schema>
```

If the databases are heterogeneous, you might use this format instead:

```
<src_catalog>.<src_schema>=<target_schema>
```

To map the target database only, use the following format:

```
<src_catalog>=<target_catalog>
```

If any of the names are null, you can safely omit the name. For example, if the source database does not have a schema, to map the table <src_table> from database <src_catalog>, you can specify

```
<src_catalog>.<src_table>=<target_catalog>.  
<target_schema>.<target_table>
```

NOTE: You do not need to specify the catalog or schema name if they are not supported in the source or target database. Do not insert null for these names.

2. In the Web Console, click **TASKS > System Parameters**.
3. Click **Cartridge**.
4. Scroll down to the parameter **Schema Mapping File for Upload**. If you are performing mapping for AQS Cache, go to the **AQS Cache Options** section.
5. Enter the path and name of the schema mapping file that you created in [Create a schema mapping file in a location that is accessible to your HPE Structured Data Manager instance](#). The contents

of the mapping file should adhere to one of the following formats, depending on whether you want to map:, above.

Cartridge mapping files

When using schema mapping to upload from a source to a target where the databases are heterogeneous, or you want to upload into a different schema name than the one from which you archived, you need a schema mapping file. A schema mapping file is useful when one database supports catalog and the other does not. You can also use the mapping file to perform explicit column type mapping.

Use the mapping file to map the following from source to target:

NOTE: These are case sensitive.

- catalog name
- schema name
- table name
- column name
- column type name
- map index name
- primary key index name

NOTE: If the table, schema, or column name contains a dot (.), the name must be enclosed in double quotes ("). If the name already contains double quotes, add one double quote for each double quote. For example, if a catalog is C, the schema name is my.schema, and the table name is new"customer", the fully qualified name is C."my.schema".new""customer"".

Syntax

The following is the syntax to use for mapping objects. You need to create a mapping for each object that you want.

```
[source_catalog_name].[source_schema_name].  
[source_table_name].[source_column_name].  
[source_column_type_name]=[target_catalog_name].  
[target_schema_name].[target_table_name].  
[target_column_name].[target_column_type_name]
```

Schema mapping supports the following:

- Column type properties on target:
- length
- precision
- scale
- Precision and scale for all numeric types.
- Only character length for all character types.

You can set these properties on any appropriate target column type. For example, you can set the length on VARCHAR2:

```
DEMARC60.dbo.ORDERHEADER.COL1.date=  
DEMARC.ORDER_HDR.COL1_1.VARCHAR2[length:18]  
DEMARC60.dbo.ORDERHEADER.COL1.money=  
DEMARC.ORDER_HDR.COL1_1.NUMBER[precision:18,scale:3]
```

NOTE: The standard method of specifying these properties is not supported because some data types contain parenthesis. For example, varchar(max). Properties should be specified within square braces as name value pairs separated by a colon.

Each line the your schema map must map only source to target mapping. For example, the following maps the data column type on the source to a VARCHAR2 type on the target.

```
DEMARC60.dbo.ORDER_HEADER.COL1.date=DEMARC.ORAER_HDR.  
COL1_1.VARCHAR2
```

The preceding example does not map any other objects, such as catalog, schema, table, or column name.

Map schemas by database

When uploading data between heterogeneous databases, it is often best to create mappings at the database level rather than the cartridge level. You must consider carefully the mapping of data types between databases and perhaps even apply some conditional logic to determine the correct mapping. This sort of data type mapping is handled through a database mapping file. If you are uploading from or to a database that is non-native to HPE Structured Data Manager, you need to create a mapping file. For native databases, you should be able to upload data from one database to another without updating the mapping file.

Non-native to native database

When uploading archiving data from a non-native database to a native database, a mapping file already exists for the target database. During the upload process HPE Structured Data Manager attempts to map the column type with the help of the existing mapping file. It also throws an exception if it finds any invalid mapping, such as the target column is incorrectly mapped, no mapping is found, or multiple mappings are found.

If the upload process fails an exception is thrown and recorded in the log file. To fix the mapping issues, edit the mapping file then rerun the upload process. For information about logging, see [Use log files, on page 71](#). For information about editing mapping files, see [Database mapping files, on the next page](#).

Native to non-native database

When uploading to a non-native database from a native database, a mapping file does not exist. In this case, HPE Structured Data Manager creates a mapping file when you run the upload process. If invalid mappings are found, the upload process fails, and exceptions are thrown and logged. To fix the mapping issues, edit the mapping file then rerun the upload process. For information about editing mapping files, see [Database mapping files, on the next page](#).

Non-native to non-native database

When uploading from a non-native database to a non-native database, a mapping file does not exist. In this case, the upload process creates a mapping file. If invalid mappings are found the upload process fails, and exceptions are thrown and logged. To fix the mapping issues, edit the mapping file then rerun the upload process. For information about editing mapping files, see [Database mapping files, below](#).

Database mapping files

When uploading from or to non-native databases, you may use the database mapping files to handle the transposition of data types between databases. HPE Structured Data Manager supports and includes mapping files for the following databases:

- DB2
- Oracle
- SQL Server
- Sybase

The mapping files are XML files that specify how one data type should be mapped to another. The mapping can be:

- Simple— One data type is mapped to another without any considerations.
- Database specific— The data type mapping depends on which database is used as the source. For example, if the data type is real, it is mapped to float for SQL Server; otherwise, it is mapped to NUMBER.
- Auto-expansion— The data types being mapped have a precision, length, or scale values but have different maximums so the mapping depends on the actual value. For example, if length exceed 4001 when mapping from varchar to a VARCHAR2 data type, it is mapped to CLOB instead of a VARCHAR2.

For the elements and structure required in a mapping file, see the mapping file schema `dataTypeMapping.xsd`. It is located in the directory:

```
<install_directory>/obt/foundation/components/config
```

where `<install_directory>` is the location where HPE Structured Data Manager is installed.

Example

The following is a example of simple mapping:

```
<mapping>
  <columnInput>real</columnInput>
  <columnOutput>NUMBER</columnOutput>
</mapping>
```

In this example, the real data type in the source is mapped to the NUMBER data type in target.

Example

The following is an example of database specific mapping:

```
<mapping>
  <columnInput>real</columnInput>
  <exceptions>
    <exception>
      <sourceDBName>
        <![CDATA[.*(sql).*(server)*.]]>
      </sourceDBName>
      <columnOutput>float</columnOutput>
    </exception>
  </exceptions>
  <columnOutput>NUMBER</columnOutput>
</mapping>
```

In this example, the real data type is mapped to NUMBER for all source databases. However, if the source is a SQL Server database, real is mapped to float.

NOTE: The upload process tries to match sourceDBName in the mapping element to the folder name generated in the /obt/config/dbms directory. For example, the regular expression ".*(sql).*(server)*." is matched with Microsoft_SQL_Server_9_0.

Example

The following is an example of auto-expansion mapping:

```
<mapping applyRule="VARCHAR2_RULE">
  <columnInput>varchar</columnInput>
  <columnOutput>VARCHAR2</columnOutput>
</mapping>
<propertiesExceptions>
  <propertiesException ruleName="VARCHAR2_RULE">
    <condition minLength="4001">
      <columnOutput>CLOB</columnOutput>
    </condition>
    <condition isLengthSet="false">
      <columnOutput length="20">VARCHAR2</columnOutput>
    </condition>
  </propertiesException>
</propertiesExceptions>
```

In this example, varchar is mapped to VARCHAR2. Also, the ruleName attribute is set to VARCHAR2_RULE and applies to the mapping. The rule specifies that if the length of varchar exceeds 4000, it is mapped to CLOB.

The propertiesException element is used to create rules. You can apply rules to mapping or exception elements. The condition elements are executed in order. When a condition is met, the data type specified in the columnOutput element is used. No subsequent conditions are applied. The condition element supports the following attributes:

- maxLength
- minLength
- maxPrecision
- minPrecision

- maxScale
- minScale
- isLengthSet
- isPrecisionSet
- isScaleSet

Create and edit database mapping files

If your source and target databases are databases supported by HPE Structured Data Manager, you do not need to edit or create a mapping file because a mapping file is already provided. However, if a database other than those supported is involved, you will need to create and edit a mapping file.

To create a mapping file

1. Open a command window and navigate to the directory `<install_directory>/obt/bin`, where `<install_directory>` is the directory where the software is installed.

2. Run the script `create_datatype_map`.

On Windows:

```
create_datatype_map.bat
```

On UNIX:

```
create_datatype_map.sh
```

The script connects to target database to create mapping for the source to target column type.

NOTE: Using the `create_database_map` is optional. However, running the script allows you to identify issues with data type mapping before running the upload process.

3. In a text editor, open the mapping file created by the script and fill in the data types for the target column. Skeleton mapping elements are added to the end of the mapping file.

When editing the mapping file, consider the following:

- The schema `dataTypeMapping.xsd` provides the usage information for
 - auto-expanding a target data type based on certain properties.
 - mapping for a specific source database is supported.
 - setting certain target column data type properties.
- The `sourceDBname` and `columnInput` elements can accept a regular expression. Regular expressions should be inside a CDATA section.

Validate the database mapping file

To validate your mapping file against the schema (`dataTypeMapping.xsd`), run the following command:

On Windows:

```
create_datatype_map.bat <mapping_file_name> -v
```

On UNIX:

```
create_datatype_map.sh <mapping_file_name> -v
```

where <mapping_file_name> is the name of your mapping file.

Create mapping by connecting to the source database

To create data type mapping in your mapping file by connecting to the source database, run the following command:

On Windows:

```
create_datatype_map.bat -f
```

On UNIX:

```
create_datatype_map.sh -f
```

The script connects to the source and target database to create mapping for the source to target column type.

Extracting binary zero in Oracle CHAR or VARCHAR

Some Oracle instance may have binary zero (CHR(0)) in the character strings. Columns containing CHR(0) characters can be extracted by SDM but fails refreshing the AQS cache.

In dbvisualizer the data would appear as:

| * | A | B | C |
|---|-------|----|------|
| 1 | a□□□□ | b□ | □□□□ |
| 2 | a□□□□ | b□ | □□□□ |
| 3 | a□□□□ | b□ | □□□□ |

1. Create a SDM masking function

In SDM home directory \${home}\extensions\runtime\masking

Create a groovy file: cleanzerobin.groovy with the following content.

```
Object mask(Object _in) {  
    return _in.replaceAll("[\\x00-\\x00]", " ");  
}  
  
Object revert_mask(Object _in){  
    return _in;  
}
```

2. Add cleanzerobin masking function.
3. Apply the cleanzerobin masking function to dirty columns.
4. Run the business flow with AQS cache enabled and you should have all the binary zeros replaced with space.

5. In dbvisualizer the data in AQS_CACHE now appear as:

| * | A | B | C | JOB_RUN_ID |
|---|---|---|---|------------|
| 1 | a | b | | 375 |
| 2 | a | b | | 375 |
| 3 | a | b | | 375 |

Chapter 14: Perform history schema reconciliation

When using database to database archiving, you need to ensure that your target (history) schema remains synchronized with your source database schema. If the managed tables in the source database have been modified in some way (for example, columns added or deleted), you need to reconcile the history schema where data was archived to match those changes.

This chapter describes how to perform the history schema reconciliation to update your archived data.

You must run the history schema reconciliation only when you have changes in your source database that are not reflected in the history (target) database. For example, if some columns in the source tables have been added or deleted, you would then need to run history schema reconciliation.

- [About the history schema reconciliation process, below](#)
- [Overview of the history schema reconciliation process, on page 124](#)
- [Prepare to run the history schema reconciliation, on page 124](#)
- [Discover differences between the source and target databases, on page 127](#)
- [Generate the XML script, on page 128](#)
- [Edit the XML script, on page 129](#)
- [Deploy the XML script, on page 132](#)
- [Verify results, on page 133](#)
- [Check the log files, on page 133](#)

About the history schema reconciliation process

The history schema reconciliation process allows you to keep your source and target database schemas consistent. Over time, you inevitably make changes to your source database objects, such as adding new columns or upgrading versions. For example, if the US-based ACME company expanded to an overseas market, they would need to add a currency-type (dollars, euros, yen) column to their Orders table. They would need to add this column to all the tables and back-fill the old orders with USD (United States Dollars).

To reflect these types of changes in your target, the history schema reconciliation process discovers changes to the managed tables of your source schema and generates a corresponding XML script for your target schema.

You should perform a history schema reconciliation after you complete any of the following activities:

- Upgrade the source database.
- Apply an application patch to the source database that modifies a managed table.
- Customize the database resulting in changes affecting the managed tables of the cartridges.

When you perform the history schema reconciliation process, it performs the following tasks for you:

- Determines if there are differences in the managed tables of the source and target schema by comparing attributes:
 - column datatype
 - column datatype size
 - column nullable status
 - column default value
 - missing individual columns
 - dropped columns
 - added columns
 - sets of columns attributes
- Detects indexes that were modified in the source and indexes that exist in the target but do not exist in the source and drops them. New indexes created in the source will not be detected. However, if you want to add those indexes into the history schema, you can create them:
- In history reconciliation by modifying the `<prefix>_build.xml` file, which is located in the home directory under `/businessflow/schema_upgrade/`, before you deploy it. For more information about the home directory, see [Configure the Web Console, on page 17](#).
- In an archive run.
- By adding the indexes manually directly to the database on the history side.

See [About the history schema reconciliation process, on the previous page](#).

NOTE: Tables in the target database should not have more than 20 indexes due to the limitation with Oracle on using union views over database links.

If there are some indexes that you do not want to remove, modify the `<prefix>_build.xml` file, which is located in the home directory under `/businessflow/schema_upgrade/`, to delete the respective DROP INDEX statements before you deploy the script.

- Generates an XML script in the home directory under `/businessflow/schema_upgrade` to synchronize the source and target databases. The script can add, drop, or modify columns, and automatically populate new columns, if the new columns are not set to NULL. For example, suppose you change a LEDGER table to insert the AMOUNT_PLASTIC column, which tracks the amount of plastic in every bottle. The AMOUNT_PLASTIC column must have a numerical value because every bottle has some plastic in it. The XML script automatically fills this non-null column with a zero and flags it for editing.

The XML script contains searchable markers that you need to edit to make the current source and target managed tables match. After you have edited the XML script, you can deploy it in test, integration, and production environments. HPE recommends carefully testing the XML script in each environment before progressing to the next.

By adding the history schema reconciliation process to the patch or upgrade process you can maintain consistency between the source and target as well as in the managed tables used by the cartridges in your business flows. For more information on cartridges and business flows, see the *HPE Structured Data Manager Concepts Guide* or the *HPE Structured Data Manager Developer's Guide*.

Overview of the history schema reconciliation process

The history schema reconciliation process consists of the following steps, which are described in greater detail in the following sections:

1. Run the history schema reconciliation script in checkonly mode.
2. Generate the XML script.
3. Edit the XML script.
 - a. Search for TODOs.
 - b. Add SQL statements where needed.
 - c. Add data to any columns designated not-null.
 - d. Remove or add INDEX statements where needed.
4. Deploy the XML script in the test environment. See [Deploy the XML script, on page 132](#).
5. Perform database tests to determine proper stability and function.
6. Diagnose any issues.
7. Edit the XML upgrade script again.
 - a. Search for TODOs.
 - b. Add or edit SQL statements as needed.
 - c. Add or alter data to not-null columns as needed.
 - d. Refresh the target database, if necessary.
8. Repeat from [Edit the XML script., above](#) as necessary.
9. Deploy the XML script.

Prepare to run the history schema reconciliation

HPE recommends running the history schema reconciliation process in a test environment or on a copy of the source database before deploying in the integration or production environments.

Before beginning the reconciliation, ensure the tasks in these sections have been completed:

- [Verify current installation, below](#)
- [Verify completed jobs and the environment name, on the next page](#)
- [Verify necessary changes, on page 126](#)

Verify current installation

You must have the most recent version of HPE Structured Data Manager and the most recent patch installed.

Ensure that all the necessary requirements for HPE Structured Data Manager are met before you install the software.

See *HPE Structured Data Manager Installation Guide*

Verify current software version

Ensure that you are running the latest version of HPE Structured Data Manager.

1. Log on to the source server as a user with the appropriate permissions.
2. Navigate to the following directory:
`<71_installerdir>/obt/bin`
where `<71_installerdir>` is the location where you installed HPE Structured Data Manager 7.1.
3. Type the appropriate command:

| Operating System | Command syntax |
|------------------|--------------------------------|
| UNIX | <code>./version_info.sh</code> |
| Windows | <code>version_info.bat</code> |

The script displays the version of HPE Structured Data Manager that is installed.

Verify current patch number

Ensure that you have the latest patch installed for HPE Structured Data Manager.

1. Log on to the source server as a user with the appropriate permissions.
2. Navigate to the following directory:
`<71_installerdir>/obt/patches/`
3. Look at the names of the installed patch files in the `patch.install.log` to determine the highest cumulative patch number.

Example

`SDM712_0000 installed`

4. Check the HPE Software Patches search website for the most recent patch number of HPE Structured Data Manager.
5. If the patch number is lower than the most current patch available, then download the current patch and install it.

For more information on patch installation see the patch ReadMe file.

Example

`HPDBArchiving_00003.text`

Verify completed jobs and the environment name

Use the Web Console to check for running and in-process database to database archiving business flows and jobs, and to determine the name of the environment.

To determine the status of the jobs

1. Start the Web Console and login to the Web Console.
2. From the Web Console launch page, click **Business Flow Management**.
3. Click **TASKS > Monitoring > Overview** from the menu at the top of the page.
The overview page displays the status of your business flows as well as any currently running commands.

To determine that no archiving jobs are scheduled to be run:

1. Start the Web Console and login to the Web Console.
2. From the Web Console launch page, click **Business Flow Management**.
3. Click **TASKS > Scheduler**.
The Scheduler displays the list of scheduled jobs.

To determine the name of the environment

1. Start the Web Console and login to the Web Console.
2. From the Web Console launch page, click **Business Flow Management**.
3. Click **ENVIRONMENT > Manage** from the menu at the top of the page.
The Manage Environments page displays all of the environments in your installation.

Verify necessary changes

Run the history schema reconciliation script in checkonly mode to verify if changes are required in the target.

1. Open a command window.
2. Navigate to the following directory:
`<71_installerdir>/obt/install`
where `<71_installerdir>` is the location where you installed HPE Structured Data Manager 7.1.
3. Run the `prepare_history_upgrade` script using the appropriate syntax:

| Operating System | Command syntax |
|------------------|---|
| UNIX | <code>prepare_history_upgrade.sh -e <environment_name> [-key <EKEY>] -checkonly</code> |
| Windows | <code>prepare_history_upgrade.bat -e <environment_name> [-key <EKEY>] -checkonly</code> |

| Parameter | Description |
|---------------------------------------|--|
| <code><environment_name></code> | The name of the environment. |
| <code><EKEY></code> | The case-sensitive encryption key you defined during installation. The <code><EKEY></code> is optional. If you place the encryption key into the |

| Parameter | Description |
|------------|---|
| | command line, it is visible on the screen. If you do not use the <EKEY>, you are prompted to type the encryption key later. |
| -checkonly | The command to compare the source and target databases without generating an XML script. This option lets you know if a schema reconciliation is necessary. |

UNIX example

```
./prepare_history_upgrade.sh -e ENV1 -key obtkey -checkonly
```

TIP: To see all the options for the history schema reconciliation script, use the -help option when you run it. For example, `prepare_history_upgrade.bat -help`.

If no changes are discovered, the following message is returned:

```
OLTP and History match...
```

If schema changes are discovered, the following message is returned:

```
OLTP and History do not match. Need to create a History schema upgrade script.
```

NOTE: If the schemas for your production database and your archived historical data match, there is no need to run the history schema reconciliation script.

Discover differences between the source and target databases

Your existing change management processes should track changes from patches, customizations updates, or DML (Data Manipulation Language). The history schema reconciliation helps you get the most from the system by ensuring that the source and target database table definitions match.

Differences between the source and target databases are typically the result of patches or customizations applied to the source database tables.

To identify the scope of changes

1. Use your company's database change management process to note, record, and communicate any changes to the source database tables managed by the cartridges in your business flows.
2. Log in to the repository, and run a query similar to the following to retrieve a list of all the managed tables.

```
SELECT a.appspack_name, mt.table_catalog, mt.table_schema, mt.table_name
      FROM obtpa_appspacks a,
           obtpa_appspack_table_uses atu,
           obtpa_managed_tables mt
WHERE a.id = atu.appspack_id
      AND mt.id = atu.managed_table_id
ORDER BY 1,2,3,4;
```

3. Examine the patch upgrades for changes to the managed tables by looking at the log or the DML for the patch or upgrade.
4. Visually examine the source database to determine what updates have been applied to the managed tables.

Generate the XML script

The XML script allows you to update the target database to make it match the source database.

1. Open a command window.
2. Navigate to the directory that contains the `prepare_history_upgrade` script:
`<71_installerdir>/obt/install`
where `<71_installerdir>` is the location where you installed HPE Structured Data Manager 7.1.
3. Run the `prepare_history_upgrade` script using the appropriate syntax:

| Operating System | Command syntax |
|------------------|--|
| UNIX | <code>./prepare_history_upgrade.sh -e <environment_name> [-key <EKEY>] [-keepfiles] [-outputDir <output_dir>] [-prefix <prefix>] [-overwrite]</code> |
| Windows | <code>prepare_history_upgrade.bat -e <environment_name> [-key <EKEY>] [-keepfiles] [-outputDir <output_dir>] [-prefix <prefix>] [-overwrite]</code> |

| Parameter | Description |
|---------------------------------------|---|
| <code><environment_name></code> | The name of the environment. |
| <code><EKEY></code> | The case-sensitive encryption key you defined during installation. The <code><EKEY></code> is optional. If you place the encryption key into the command line, it is visible on the screen. If you do not use the <code><EKEY></code> command here, you are prompted to type a hidden encryption key later. |

| Parameter | Description |
|--------------|---|
| -keepfiles | Retains all of the working files used by the reconciliation process. The files are placed in the same directory as the reconciliation script or the specified output directory. |
| <output_dir> | Use the -outputDir keyword to set a different output directory for the generated files. By default, the output is placed in the directory where the script is located. |
| <prefix> | <p>The name chosen to specify a unique prefix to track various versions of the generated reconciliation script. Use sequential numbers, patch numbers, release numbers, or timestamps as unique names.</p> <p>The prefix is optional. By default, the output of the reconciliation script is named: HISTORY_UPGRADE_build.xml</p> |
| -overwrite | <p>An optional command you can use to overwrite a previously generated XML script of the same name.</p> <p>You cannot overwrite a script without using the -overwrite option. If you use the same name, the script fails with an error message.</p> |

Example

```
./prepare_history_upgrade.sh -e oracle_oltp_env -outputDir  
../businessflow/schema_upgrade
```

TIP: To see all the options for the history schema reconciliation script, use the -help option when you run it. For example, `prepare_history_upgrade.bat -help`.

The generated XML script is saved to the current working directory, or the output directory you specify.

Edit the XML script

Edit the XML script as necessary to ensure that the source database matches the target database. You can add data definition language (DDL) operations, such as index or keys that do not exist in the generated XML script. You can add DML for deletes, updates, or inserts that do not exist in the generated XML script.

The generated XML script handles schema structure changes including:

- adding a column
- dropping a column
- revising the column data type
- revising the column width

If the changes in source database also involve data transformations, such as the following, then you must edit the reconciliation script to ensure the database schemas match:

- using application code to transform data
- adding a default column
- adding or filling a not null column

Data changes that require editing are marked by one of the following TODO comments:

| Type | Statement |
|---------------------|---|
| default seeding | <code><!-- TODO: Correct the following default SQL seeding statement, it is most likely to be incorrect.--></code> |
| data transformation | <code><!-- TODO: If needed, add data transformation SQL statement here by uncommenting the obtsql task and replacing the "SQL Statement" with a real SQL DML statement --></code> |

Not all of the TODO statements in the generated XML script require editing.

NOTE: If you want to use the greater than (>) or less than (<) expressions, you need to quote them according to XML 1.0 rules. The greater than (>) symbol should be written as `>`; and the less than symbol (<) should be written as `<`;

To edit the XML script

1. Navigate to the directory containing the generated XML script.

```
<app_data>/HPE/HPEOBT/businessflow/schema_upgrade/<prefix>_build.xml
```

| Parameter | Description |
|--------------------------------------|--|
| <code><71_installerdir></code> | The location where you installed HPE Structured Data Manager 7.1. |
| <code><prefix></code> | The name used to specify a unique prefix to the XML script. If you did not specify a prefix, the XML script is called HISTORY_UPGRADE_build.xml. |

2. Use an XML editor or text editor to open the XML script.
3. Save the file with a unique name.
4. Search for comments containing the text TODO.

Example

```
<!-- TODO: If needed, add data transformation SQL statement here by  
uncommenting the obtsql task and replacing the "SQL Statement" with a real SQL  
DML statement -->  
<!--  
<obtsql connectionSource="${cs.HIST}" autocommit="true" echo="true"> SQL  
statement </obtsql>  
-->
```

The TODO comments are accompanied by the obtsql element which is used to contain the optional SQL statement. The statements can appear before or after the specified update.

5. Uncomment the obtsql element that you want to edit.

Example

```
<!-- TODO: If needed, add data transformation SQL statement here by
uncommenting the obtsql task and replacing the "SQL Statement" with a real SQL
DML statement -->
<obtsql connectionSource="${cs.HIST}" autoccommit="true" echo="true"> SQL
statement </obtsql>
```

6. Replace the placeholder text with a valid SQL statement.

Example

```
<!-- TODO: If needed, add data transformation SQL statement here by
uncommenting the obtsql task and replacing the "SQL Statement" with a real SQL
DML statement -->
<obtsql connectionSource="${cs.HIST}" autoccommit="true" echo="true">
UPDATE TABLE_NAME SET CURRENCY_CODE = ' '
WHERE CURRENCY_CODE IS NULL
</obtsql>
```

NOTE: Ensure that the SQL statements do not see tables that have been removed from the schema.

You can edit the entire SQL statement or just a portion.

For example, you can change the currency code from an empty value to United States dollars (USD).

Example

```
<obtsql connectionSource="${cs.HIST}" autoccommit="true"
echo="true">
UPDATE TABLE_NAME SET CURRENCY_CODE = 'USD'
WHERE CURRENCY_CODE IS NULL
</obtsql>
```

If adding non-null columns, the following placeholders are used:

| Non-null column type | Placeholder value |
|----------------------|--|
| date | 1970 |
| number | 0 |
| string | an empty space between single quotes ' ' |

7. Edit the TODO comment with a searchable string.

Example

```
<!-- Edited by John Doe on Dec. 8, 2008 -->
<obtsql connectionSource="${cs.HIST}" autoccommit="true" echo="true"> UPDATE
TABLE_NAME SET CURRENCY_CODE = ' '
WHERE CURRENCY_CODE IS NULL </obtsql>
```

8. Save the edited XML script.

NOTE: The XML script elements that you define must be compatible with the environment where the XML script is prepared and used. For example, the test database can have columns the production database does not have.

NOTE: The history schema reconciliation uses Apache ANT to deploy the changes. If you want to add additional statements to the XML script, see the Apache ANT 1.7.0 user manual.

Deploy the XML script

After you have edited the XML script, you can deploy it against test, integration, and production environments. HPE recommends carefully testing the XML script in each environment before progressing to the next. Deployment of the history reconciliation script creates interface and relocation tables and reinstalls the business flows. After you deploy the edited script, there is no undo. You should deploy and fully test the XML script before deploying it against a production database.

To deploy the XML script

1. Open a command window.
2. Log in as a user with the appropriate permissions.
3. Navigate to `/businessflow/schema_upgrade` under the home directory. See [Configure the Web Console, on page 17](#).
4. Run the `deploy_history_upgrade` script using the following syntax:

| Operating System | Command syntax |
|------------------|--|
| UNIX | <code><71_installerdir>\obt\install\deploy_history_upgrade.sh <upgrade_script></code> |
| Windows | <code><71_installerdir>\obt\install\deploy_history_upgrade.bat <upgrade_script></code> |

| Parameter | Description |
|--------------------------------------|---|
| <code><71_installerdir></code> | The location where you installed HPE Structured Data Manager 7.1. |
| <code><upgrade_script></code> | The name you use specify a unique version of the XML script. Use this file name to track various versions of the generated XML script. For example, you could use sequential numbers, patch numbers, release numbers or timestamps as unique names. |

Example

```
/export/home/user/HPESDM/obt/install/deploy_history_upgrade.sh HISTORY_UPGRADE_build.xml
```

5. Provide the encryption key when the script prompts you for it.

After you provide the encryption key, the script runs against your database. If the deployment is successful, you receive a message stating `Deployment successful`. You can then confirm the results as per [Verify results, below](#).

Otherwise, you receive a message stating:

`Failed to deploy. Please see the log output for more information.`

In that case, you should review the log files to determine what went wrong as per [Check the log files, below](#).

Verify results

After you have deployed the edited XML script in the test, integration or production environments, follow your company procedures for running test cases. Ensure that you run reports that would reveal empty non-null columns.

To verify results in the test environment:

- check the table definitions for expected changes.
- compare the managed tables for expected changes.

If empty columns are discovered after running test cases then you can re-seed the non-null columns. The empty columns occur when either:

- the XML script did not update all records.
- not all values were applied as expected.

Search for the SQL statements associated with non-null columns and correct the default value or SQL statement.

Check the log files

All the schema changes applied to the target database are kept in a tracking log. The log file shows the succession of steps the history reconciliation schema completed. Each one automatically verifies that the previous step completed successfully.

- [Check the .out file, below](#)
- [Check the obt.log file, on the next page](#)
- [Edit the logging properties, on the next page](#)

Check the .out file

By default, the .out file is saved to the schema reconciliation directory. If you have not specified a unique prefix, the default name is `HISTORY_UPGRADE_build.out`. The output from each run of the history schema reconciliation is appended to the .out file with a timestamp.

To check the .out file:

1. Open a command window.
2. Navigate to the schema reconciliation directory:

```
<71_installerdir>/obt/businessflow/schema_upgrade
```

where <71_installerdir> is the location where you installed HPE Structured Data Manager 7.1.

3. Open the log file.
4. Use the log files to find information on the following:
 - incorrect encryption key
 - incorrect password
 - incomplete test environment
 - incorrect copying of the code tree
5. Close the log file.

Check the obt.log file

By default, the logging level is set to INFO.

1. Open a command window.
2. Navigate to the log file:

```
<71_installerdir>/obt/log/
```

where <71_installerdir> is the location where you installed HPE Structured Data Manager 7.1.
3. Open the obt.log file.
4. Use the log files to find information on the following:
 - incorrect encryption key
 - incorrect password
 - unknown environment
 - incomplete test environment
 - out of resources, especially when applying updates
 - business flows copied to wrong location
 - incorrect copying of the code tree
5. Close the log file.

Edit the logging properties

After you have installed the product, you can edit the log4j.properties file to change the logging properties. For example, you can change:

- where the log files are kept.
- what information is logged.
- the maximum size of the log file.
- how much logging information is kept.

By default, the obt.log file is limited to 10 MB. When it exceeds the default limit, the obt.log file is renamed to a backup file, and a new obt.log file is created. By default, a total of nine backup files are kept, limiting the log files to a 100 MB maximum size.

To edit the log4j.properties file

1. Navigate to the directory that contains the log4j.properties file.
 <71_installerdir>/obt/config/
 where <71_installerdir> is the location where you installed HPE Structured Data Manager 7.1.
2. Open the log4j.properties file by using a text editor and edit appropriately.
3. Save the log4j.properties file.

The changes are applied automatically.

Chapter 15: Use AQS to access database to file output

After your data has been archived to XML, CSV, or JSON files, you can use the Archive Query Server (AQS) cache to directly access your data using SQL reporting and development tools on Windows and UNIX, and join the data with existing data in a database.

AQS uses PostgreSQL to cache your archived data for query access.

Configure the AQS cache

To use the AQS cache, you must specify a location for AQS_CACHE and configure the cache parameters in your environment:

1. If the Web Console is not open in your browser, invoke it by accessing its URL, for example, <http://localhost:8080/WebConsole>.
2. Login as the admin user, whose password you set when installing the repository, or as another user with administrative privileges.
3. Click **Business Flow Management** on the Web Console home page.
4. Choose **Environment > Active Environment > Locations**.

TIP: If necessary, you can activate a different environment by clicking **Environment > Manage**, selecting the desired environment, and clicking **Set Active**.

5. Click **Locations** from the left navigation pane.
6. Click **AQS_Cache**.
7. Enter the connection information for a PostgreSQL database instance. If you installed an instance as part of HPE Structured Data Manager installation, use that one. If not, you will need to use an existing instance or install one.

NOTE: Some PostgreSQL requirements differ from other databases. For example, PostgreSQL only allows schema, table, or column names to be a maximum of 63 characters whereas DB2 allows names up to 128 characters. Hence, if you are archiving from a source database of DB2, you need to implement schema mapping for your AQS cache to ensure that your schema, table, and column names are handled appropriately. See [Map schemas for upload or AQS cache, on page 113](#) for more information about schema mapping.

8. Click **Update**.

TIP: Notice that this page also has buttons for refreshing and clearing the cache. When you need to perform those functions, remember to return here.

9. Choose **TASKS > System Parameters > Database to File**.
10. Scroll to the **AQS Cache Options** section.
11. Set **Enable AQS Cache** to **true**.

12. Optionally, set the other AQS parameters. Click the information icons to get more information about the purpose of the parameter. For more information on setting these parameters, see *HPE Structured Data Manager Runtime Guide*.

Since you have already run your archive, you must refresh the cache to populate it with all of your archived records.

NOTE: You can specify the schema for the AQS cache in `outerbay.properties`, which is stored in the home directory under `config`. See [Configure the Web Console, on page 17](#)

13. Click **Locations** from the left navigation pane.
14. Click **AQS_Cache**.
15. Click **Refresh All Cache**.

NOTE: Refresh All Cache will not work if the archived data is in HP Records Manager as part of a Structure Records Management Solution (SRMS). In that case, you would need to migrate the archived data from SRMS.

Before migrating the SRMS data for AQS cache, ensure that you are using the same repository. Upgrade the repository from lower version to higher version, and then run **Migration**. For example, SDM version 7.21 to SDM version 7.4. Ensure that the RQS option is selected to obtain data from RM. If RQS Location is selected, you must enter the RM location.

16. The AQS cache should now be populated and you can run queries against it as you would any other database.

Part IV: Appendixes

Appendix A: Configuration and runtime parameters

The following parameters govern the running of business flows and jobs:

- [Database to file configuration parameters, below](#)
- [Database to database configuration parameters, on page 150](#)
- [File naming parameters, on page 161](#)

Database to file configuration parameters

Selecting Database to file on the Parameters page displays the configuration parameters applied to all business flows and jobs that employ database to file archiving.

The configuration parameters are divided into the following sections:

- [Core parameters, below](#)
- [Performance parameters, on page 143](#)
- [Validation parameters, on page 146](#)
- [Centera parameters, on page 147](#)
- [Amazon S3, on page 148](#)
- [Upload options , on page 149](#)
- [AQS cache options , on page 149](#)
- [HPE CM options , on page 149](#)

Core parameters

The following parameters are located in the Core section:

| Parameter name | Description |
|--------------------------------------|--|
| Allow masked data on undo and reload | By default, prevents any data that has been masked from being reloaded into the source database. Set to true if you want to reload masked data instead of the original values. |
| Compression algorithm | <p>Specifies the compression algorithm to apply to the XML and CSV files created after running the database to file archive. Valid values are No Compression and GZIP compression.</p> <ul style="list-style-type: none">• No compression—places the files in the specified directory without compression.• GZIP compression—compresses the files into GZIP format and places them in the specified directory. GZIP compression is the default. |

| Parameter name | Description |
|---|---|
| Extract file format | Specifies the file format for extracted data. Valid values are: <ul style="list-style-type: none"> • XML denormalized extracts data as XML. • CSV normalized extracts data as comma separated values. • JSON denormalized extracts data as JavaScript Object Notation. The default format is CSV normalized . |
| Job engine SQL tracing enabled | Turns on SQL tracing for the job engine repository connection. |
| Preserve temporary files | Saves the temporary files generated when running the job. |
| Primary key index location | Displays the location where the primary key indices are stored. |
| Source database location | Specifies the name of the active database. |
| Storage Retention | Is the retention policy to submit to the storage subsystem when files are stored to the storage system. Enter the number of days to retain the extracted data. If left blank, the retention behavior defaults to that of the storage system. |
| Unmask data on upload | Set to true to have data unmasked during upload from file to database. The data is unmasked if the mask used during archiving is reversible. If the mask is not reversible, no attempt is made to revert values, and masked values are uploaded as is from the file archive. |
| Use Data Movement Key to move data | Set to true to use a data movement key when required to identify rows. For Oracle databases, you can set this parameter to false to use ROWID rather than the data movement key. For non-Oracle databases, this parameter has no effect. Because database reorganization may change ROWID values, you should only consider using ROWID when you are certain that it will not change during archive. For more information about ROWID, see your Oracle documentation. |
| User index location | Displays the location where the user indexes are stored. |
| Verify Row Counts | Set to true to perform verification of row counts between the current job and its corresponding selection job. Set to false to bypass verification. True is the default value. This parameter applies only to database to database. |
| Write XSD and Summary even when there are no data files | Set to true to write XSD and summary files even when there are no data files. The default is true. |

Core parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the core parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|--------------------------------------|--|---|------------|--------------------|
| Allow masked data on undo and reload | ARCHIVE_UNDO_SELECTION RELOAD_SELECTION | ARCHIVE_UNDO_SELECTION - Snapshot qualifying HIST-Table rows into Undo-Selection Tables, RELOAD_SELECTION -Snapshot qualifying HIST-Table rows into Reload-Selection Tables | Task | BF / Tasks |
| Compression algorithm | COPY_DB_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD, WRITE_SUMMARY | Job / Task | |
| Extract file format | COPY_DB_TO_BE (D2F) | SEED_COPY_DB_TO_CSV, SEED_CSV2XML | Job / Task | |
| Job engine SQL tracing enabled | Starts trace for JobEngine Module only. All jobs | | Job | BF |
| Preserve temporary files | COPY_DB_TO_BE (D2F), UPLOAD_BE_TO_DB | SEED_COPY_DB_TO_CSV, SEED_CSV2XML, CSV2XML_CLEANUP, SEED_XML_TO_CSV | Job / Task | |
| Primary key index location | COPY_DB_TO_BE (D2F), DELETE_FROM_DB (D2F) | | Job / Task | |
| Source database location | All jobs requiring source database | | | |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|------------------------------------|---|--|------------|--------------------|
| Storage Retention | COPY_DB_TO_BE (D2F), COPY_BE_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD, WRITE_SUMMARY | Job / Task | |
| Unmask data on upload | UPLOAD_BE_TO_DB | UPLOAD_BE_TO_DB | Job / Task | |
| Use Data Movement Key to move data | Selection and Archive Jobs | During prepare it uses this parameter to generate selection views. For D2D it uses to decide if copy should populate obt_saved_rowid on history table. | Job | BF |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|---|-------------------------|---|------------|--------------------|
| User index location | COPY_DB_TO_BE (D2F) | EXTRACT_USER_INDEXES | | |
| Verify Row Counts | Selections Jobs only | UPDATE_ROWCOUNT_SUMMARY- Update rowcount summary for the run | | |
| Write XSD and Summary even when there are no data files | COPY_DB_TO_BE (D2F) | WRITE_ATT_GROUP_XSD, WRITE_SUMMARY | Job / Task | |

The table below lists the Legend Key and the Legend Value.

Table legend

Legend key: Short name of the task.

Legend value: Description of the task.

| Legend Key | Legend Value |
|------------------------|--|
| CREATE_TABLE_INDEXES | Create Indexes on destination tables |
| CSV2XML | Legend Value: File format conversion |
| CSV2XML_CLEANUP | File format conversion clean up |
| DELETE_FROM_DB | Delete archived data from database |
| EXTRACT_DELETE_INDEXES | Extract delete-index data matching selection tables to delete-index tables |
| EXTRACT_USER_INDEXES | Extract user-index data matching selection tables to user-index tables |
| SEED_COPY_DB_TO_CSV | Seed JE threads to copy Indexes and to CSV |
| SEED_CREATE_INDEXES | Seed JE Threads for creating indexes on destination tables |
| SEED_CSV2XML | Seed threads for file format conversion |
| SEED_XML_TO_CSV | Seed JE Threads for Converting XML Data to CSV format. |

Performance parameters

The following parameters are located in the performance section:

| Parameter name | Description |
|-------------------------------------|---|
| Combined delete count | Defines the maximum number of delete statements that will be executed in a single request to the database. |
| Combined statement count | Defines the maximum number of database statements which can be combined in a single query. |
| Data movement Batch size | Controls the number of driving table rows per transaction. This parameter is used for data movement operations on related parent and child rows in the same transaction, which includes database to XML file operations. The total number of rows operated on can be much larger than the value entered, and depends on the characteristics of the data. A value of zero disables intermediate commits. |
| Eligibility Analytics Configuration | <p>By default, eligibility analytics is disabled. If you want to enable eligibility analytics for a business flow, you must enable it before you run the business flow. Enabling eligibility analytics allows querying of the analytics tables for information on record eligibility.</p> <ul style="list-style-type: none"> • Select Disable the eligibility analytics to improve performance • Select Enable the eligibility analytics to enable eligibility analytics. <p>NOTE: If you select this option, you must include an interrupt step in your business flow to view the data.</p> |
| Maximum number of parallel workers | Defines the default maximum number of job workers for tasks that can take advantage of parallelism. |
| Populate Index Tables | Specifies whether the index tables should be populated. If you plan to delete your data rather than copy it, the value must be set to true. |
| Populate User-Index tables | Specifies whether the user index tables should be populated. If you plan to delete your data rather than copy it, the value must be set to true. |
| Selection batch size | Defines the number of driving table rows per transaction. This is used for selection operations that select related parent and child row IDs from the source database into the selection tables. |
| Units of work | Defines the number of units amongst which the total amount of work will be divided. Each worker picks up a whole unit at a time to ensure clear progress indication and manage the total work in units for the job engine. |

Performance parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the performance parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|-------------------------------------|--|---|------------|--------------------|
| Combined delete count | | DeleteByArchivedDataFrom DatabaseTask | Job | Tasks |
| Combined statement count | COPY_BE_TO_DB, UPLOAD_BE_TO_DB | COPY_BE_TO_DB, UPLOAD_BE_TO_DB | Job / Task | |
| Data movement Batch size | COPY_DB_TO_BE (D2F) | SEED_COPY_DB_TO_CSV | | |
| Eligibility Analytics Configuration | Selections Jobs only | | | |
| Maximum number of parallel workers | COPY_DB_TO_BE (D2F), DELETE_FROM_DB, UPLOAD_BE_TO_DB | COPY_DB_TO_CSV, CSV2XML, DELETE_FROM_DB, COPY_CSV_TO_DB | Job / Task | |
| Populate Index Tables | COPY_DB_TO_BE (D2F) | EXTRACT_DELETE_INDEXES | Job / Task | |
| Populate User-Index tables | COPY_DB_TO_BE (D2F) | EXTRACT_USER_INDEXES | Job / Task | |
| Selection batch size | Selections Jobs only | | | |
| Units of work | | | | |

The table below lists the Legend Key and the Legend Value.

Table legend

Legend key: Short name of the task.

Legend value: Description of the task.

| Legend Key | Legend Value |
|-------------------------|--|
| VALIDATE_DATA_UNCHANGED | Validate data unchanged since archival |
| WRITE_ATT_GROUP_XSD | Write XSD file for archived files |

| Legend Key | Legend Value |
|---------------|---|
| WRITE_SUMMARY | Write summary XML and XSD files for archive |

Validation parameters

The following parameters are located in the validation section:

| Parameter name | Description |
|--|---|
| Check archive consistency | Indicates whether to validate that the archive as a whole is consistent. |
| Check cardinality constraints | Indicates whether to validate that the extracted data does not violate cardinality constraints in the model instance definition. |
| Checksum algorithm | Indicates whether to run the checksum algorithm on created files. |
| Data unchanged validation | Indicates whether to validate that the data to be deleted has not changed on the source database. |
| Delete Count Validation | Indicates whether to verify the number of rows to be deleted against the expected number of rows. |
| Match rowcounts | Indicates whether to verify that rowcounts in the XML files match those in the database. All database to file jobs use this parameter to match row counts for the current one with that of the previous per table use. |
| Verify Cartridge version for DB deletion | Indicates whether to verify that the AppsPack version used during database deletion is the same as the one used during database extraction. |
| XML Schema validation | <div>Indicates whether to validate that the XML files do not violate their XML schema. TIP: For database to file archiving, if you have large BLOB, CLOB, LONG, or LONG RAW fields, HPE recommends leaving the value of the XML schema validation parameter set to false.</div> |
| XML file checksum validation | Indicates whether to validate that the XML file checksums have not changed. |

Validation parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the validation parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

| Parameter name | Job | Task | Recovery |
|--|--|---|------------|
| Check archive consistency | All D2F and D2F-Lite Jobs | VALIDATE_ARCHIVE_CONSISTENCY- Verify XML archive is self-consistent | Job / Task |
| Check cardinality constraints | COPY_DB_TO_BE (D2F) | CSV2XML | Job / Task |
| Checksum algorithm | All D2F and D2F-Lite Jobs (Except DELETE-Jobs) | VALIDATE_CHECKSUMS - Verify checksums for files | Job / Task |
| Data unchanged validation | DELETE_FROM_DB (D2F) | SEED_VALIDATE_DATA_UNCHANGED, VALIDATE_DATA_UNCHANGED | Job / Task |
| Delete Count Validation | DELETE_FROM_DB (D2F) | DELETE_FROM_DB | Job / Task |
| Match rowcounts | DELETE_FROM_DB (D2F) | UPD_ROWCOUNT_SUMMARY | Job / Task |
| Verify Cartridge version for DB deletion | | | |
| XML Schema validation | All D2F and D2F-Lite Jobs (Except DELETE-Jobs) | VALIDATE_FILE_CONTENTS - Validate contents of XML files | Job / Task |
| XML file checksum validation | All D2F and D2F-Lite Jobs (Except DELETE-Jobs) | VALIDATE_CHECKSUMS - Verify checksums for files | Job / Task |

Centera parameters

The following parameters are located in the Centera section:

| Parameter name | Description |
|-------------------------------|---|
| Minimum retention duration | Defines the minimum number of days that data must be retained before it can be deleted. -1 means use the default. |
| Minimum retention policy name | The name of the retention policy to use. |

Centera parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the Centera parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|-------------------------------|--|---|----------|--------------------|
| Minimum retention period | D2F and D2F Lite copy jobs, DB_TO_BE_DIRECT, COPY_DB_TO_BE and COPY_BE_TO_BE | Used by Various tasks in D2F: FilePermissionValidationTask, CopyStorageToStorageStartTask, LWArchiveExtractDescriptorTask, ArchiveExtractDescriptorTask. On WebConsole: 1) FILE_PERMISSION_CHECK - Check File And Directory Permissions of Copy by selection from database to backend. 2) Copy from backend to backend startup | Task | Task |
| Minimum retention policy name | D2F and D2F Lite copy jobs, DB_TO_BE_DIRECT, COPY_DB_TO_BE and COPY_BE_TO_BE | Used by Various tasks in D2F: FilePermissionValidationTask, CopyStorageToStorageStartTask, LWArchiveExtractDescriptorTask, ArchiveExtractDescriptorTask. On WebConsole: 1) FILE_PERMISSION_CHECK - Check File And Directory Permissions of Copy by selection from database to backend. 2) Copy from backend to backend startup | Task | Task |

Amazon S3

The following parameters are located in the Amazon S3 section:

| Parameter name | Description |
|---------------------------------|--|
| Maximum number of read attempts | Is the number of times you want to try to read a file from Amazon S3 if an exception occurs during the read operation. |

Upload options

The following parameters are located in the Upload Options section:

| Parameter name | Description |
|----------------------|---|
| Enable Add Column | Set to true to add a target column during upload for any source column that does not already have a corresponding target column. Set to false if you do not want to add such columns during upload. |
| Enable Drop Column | Set to true to drop a target column during upload if it does not exist in the source table. Set to false if you do not want to drop such columns during upload. |
| Enable Modify Column | Set to true to modify the target column during upload if it is different from the source column. Set to false if you do not want to change the target column during upload. |

AQS cache options

The following parameters are located in the AQS Cache Options section:

| Parameter name | Description |
|---|---|
| Drop AQS cache schema and tables after clearing the cache | Set to true to drop the AQS cache tables after clearing the cache, if no data is found. If no tables are found, drop the schema as well. |
| Enable AQS cache | Set to true to create the AQS cache at the end of a database to file archive. Set to false if you do not want an AQS cache. |
| Enable add column to cache tables | Set to true to add a cache column during archive for any source column that does not already have a corresponding cache column. Set to false if you do not want to add such columns during archive. |
| Enable de duplication check against cache | Set to true to check for duplicate entries in the cache. Set to false to prevent de duplication from running. |
| Object mapping file for AQS cache | Specifies a mapping file to use when creating the cache. |
| Target schema name | Choose from the list of values what you want to use as the schema name. |

HPE CM options

NOTE:

- `HPRMAdditionalField.xml` file can be used to set the Parameter value from HPE Structured Data Manager.

- `HPRMAdditionalField.xml` is used to set the Parameter value if the parameter is already defined.
- Data type `Text` in RM is mapped to data type `String` in HPE Structured Data Manager.
- Property **LEVEL** can be added in `HPRMAdditionalField.xml` file which allows the metadata to go to either container or document or both.
- The values for **LEVEL** can be 'container', 'document' or 'both'. If **LEVEL** is not defined then we assume the value is 'both'.

The following parameters are located in the HPE CM Options section:

| Parameter name | Description |
|---|---|
| Additional fields configuration file name | Specifies the name of the file containing the additional HPE Content Manager field definitions. |
| Additional fields profile name | Specifies the name of the HPE Content Manager profile. |
| Classification | Specifies the HPE Content Manager classification name or UID. NOTE: For example, if the UID is 26 then the format is ID:26. |
| Container record type | Specifies the HPE Content Manager container record type name or UID. |
| Document Record type | Specifies the HPE Content Manager document record type name or UID. |
| External Reference | Specifies the HPE Content Manager external reference. |
| Notes | Specifies the HPE Content Manager note. |
| Retention schedule | Specifies the HPE Content Manager retention schedule. NOTE: : Enter the parameter as Schedule Number instead of Schedule Title. |

Database to database configuration parameters

Selecting database to database on the Parameters page displays the configuration parameters applied to all business flows and jobs that employ database to database archiving.

The configuration parameters are divided into the following sections:

- [Core parameters, on the next page](#)
- [Performance parameters, on page 152](#)
- [Advanced selection parameters, on page 158](#)
- [Distributed instance parameters, on page 159](#)

- [Archive access parameters](#) , on page 156
- [File naming parameters](#), on page 161

Core parameters

The following parameters are located in the core section:

| Parameter name | Description |
|--------------------------------------|---|
| Allow masked data on undo and reload | By default, prevents any data that has been masked from being reloaded into the source database. Set to true if you want to reload masked data instead of the original values. |
| Job Engine SQL tracing enabled | Turns on SQL tracing for the job engine repository connection. |
| Use Data Movement Key to move data | Set to true to use a data movement key when required to identify rows. For Oracle databases, you can set this parameter to false to use ROWID rather than the data movement key. For non-Oracle databases, this parameter has no effect. Because database reorganization may change ROWID values, you should only consider using ROWID when you are certain that it will not change during archive. For more information about ROWID, see your Oracle documentation. |
| Verify Row Counts | Set to True to perform verification of row counts between the current job and its corresponding selection job. Set to False to bypass verification. True is the default value. |

Core parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the Core parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|----------------|---------------|---|----------|--------------------|
| Allow masked | ARCHIVE_UNDO_ | ARCHIVE_UNDO_SELECTION - Snapshot qualifying HIST-Table rows into | Task | Tasks |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|------------------------------------|---|---|----------|--------------------|
| data on undo and reload | SELECTION RELOAD_ SELECTION | Undo-Selection Tables, RELOAD_ SELECTION -Snapshot qualifying HIST- Table rows into Reload-Selection Tables | | |
| Job Engine SQL tracing enabled | Starts trace for JobEngine Module only. All jobs | ARCHIVE_UNDO_SELECTION - Snapshot qualifying HIST-Table rows into Undo-Selection Tables, RELOAD_ SELECTION -Snapshot qualifying HIST- Table rows into Reload-Selection Tables | Job | BF |
| Use Data Movement Key to move data | Selection and Archive Jobs | During prepare it uses this parameter to generate selection views. For D2D it uses to decide if copy should populate obt_ saved_rowid on history table. | Job | BF |
| Verify Row Counts | Selection and Archive Jobs | During prepare it uses this parameter to generate selection views. For D2D it uses to decide if copy should populate obt_ saved_rowid on history table. | Task | Tasks |

Performance parameters

The following parameters are located in the performance section:

| Parameter name | Description |
|----------------------------|---|
| Analyze Tables | Indicates whether tables must be analyzed and table-statistics gathered. This parameter does not apply to advanced selection. |
| Compress tables on History | Compresses the history tables when set to True and is only used during first time deployment. This parameter is applicable only to Oracle. |
| Data movement Batch size | Controls the number of driving table rows per transaction. This parameter is used for data movement operations on related parent and child rows in the same transaction, which includes database to XML file operations. The total number of rows operated on can be much larger than the value entered, and depends on the characteristics of the data. A value of zero disables intermediate commits. |
| Data movement method | Defines the method to use to move the data. <ul style="list-style-type: none"> Fully transactional—Related parent and child rows are copied and deleted in the same transaction. Referential integrity on history is maintained at all times. |

| Parameter name | Description |
|--|---|
| | <ul style="list-style-type: none"> Table-Parallel—Related parent and child rows are copied in separate transactions per table. Deletes happen in a separate transaction, but related parent and child rows are deleted in the same transaction. This option allows faster performance, but archive access will temporarily have inconsistent data. <p>This parameter does not apply to advanced selection.</p> |
| Eligibility Analytics Configuration | <p>By default, eligibility analytics is disabled. If you want to enable eligibility analytics for a business flow, you must enable it before you run the business flow. Enabling eligibility analytics allows querying of the analytics tables for information on record eligibility.</p> <ul style="list-style-type: none"> Select Disable the eligibility analytics to improve performance Select Enable the eligibility analytics to enable eligibility analytics. <p>NOTE: If you select this option, you must include an interrupt step in your business flow to view the data.</p> |
| Estimate Percentage | Defines the estimate percentage for gathering statistics. This value is passed to DBMS_STATS.GATHER_TABLE_STATS to compute statistics. Enter a number between 0 and 100. This parameter does not apply to advanced selection. |
| History Index Tablespace Name | (Oracle only) Specifies a separate tablespace for storing indexes. |
| Maximum number of parallel workers | Defines the number of tasks that can be executed concurrently in an action. This parameter only affects jobs for the product and does not alter the database vendor's parallelism. If the server is exclusively used, the suggestion is 2 times the number of server CPUs. This parameter does not apply to advanced selection. |
| Minimum driving ATT chain size to process in a single commit for Fully Transactional operation | The driving ATT chain size greater than or equal to this value will be processed in a single commit for Fully Transactional (FT) operation. Only set this parameter when Advance Selection is executed and the driving ATT may be highly skewed because of ATT chaining. Entering the value 0 disables the feature. |
| Number of driving table rows per commit for Fully Transactional operation | Defines the maximum number of driving table rows per commit for fully transactional operation. A driving table row includes all related child rows. The total number of rows operated on can be much larger than the value entered, and will depend on the characteristics of the data. A value of zero disables intermediate commits. Also controls the commit size for table parallel delete operation, driven by driving table rows. |

| Parameter name | Description |
|---|--|
| Number of rows per commit for Table Parallel copy operation | Controls the transaction size for the copy operation for table parallel data movement's copy operation. That is, it will not commit rows greater than the value set for this parameter. A value of zero disables intermediate commits. |
| Number of tasks to be created for Data Movement | Defines the maximum number of tasks (chunks of work) that will be created for data movement. More than one task allows parallel execution if number of workers is set greater than 1. For table parallel copy task, if the value is less than the total number of tables in a cartridge, then one task will be created per table. |
| Parallel degree for creating History indexes | Specifies the parallel degree to use for index creation on history. |
| Selection batch size | Defines the number of driving table rows per transaction. This is used for selection operations that select related parent and child row IDs from the source database into the selection tables. |
| Type of Snapshots for Lookup tables | <p>This parameter controls the type of snapshot to be created on history for lookup tables. The possible types are:</p> <ul style="list-style-type: none"> • SNAPSHOT_WITH_HISTORY—Retain the row history for snapshots from previous runs. Eligibility analytics is not available for lookup tables if you select this option. • SNAPSHOT_WITHOUT_HISTORY—Do not retain any row history for snapshots from previous runs. <p>NOTE: Snapshots for lookup tables is not supported for bulk archiving.</p> <p>NOTE: When setting this parameter, make sure that you have unique keys on the lookup tables.</p> |
| Type of indexes to create on History | This parameter controls the type of indexes to re-create on History after archiving the data into History. Select Operate on All to improve the performance of access to your archived data. This parameter does not apply to advanced selection. |

| Parameter name | Description |
|------------------------------------|--|
| Type of indexes to drop on History | <p>This parameter controls the type of indexes to drop on History before archiving the data into History. This parameter does not apply to advanced selection.</p> <p>Dropping indexes can improve the performance for copying data, but the indexes will need to be recreated after the archive is completed. If the quantity of data to be copied is small compared to the quantity of already archived data, this could impact the overall archive performance. Discuss with your database administrator (DBA) to determine whether it is better to drop and recreate the indexes or leave them there for the copy.</p> |

Performance parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the performance parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|--|--------------------------------|--|----------|--------------------|
| Analyze tables | Archive Job | Determining whether to gather statistics on History tables | Task | Tasks |
| Compress tables on History (only used during first time deployment). | | | | |
| Data movement Batch size | | | | |
| Data movement method | Archive Job | PARALLELIZE_MOVE_O2H/H2O - Parallelize moving data from source to history and vice versa | Task | BF / Tasks |
| Eligibility Analytics Configuration | Standard and PDM Selection Job | ARCHIVE_ANALYZE_ELIGIBILITY - Perform eligibility analytics on snapshot rows | Task | BF / Tasks |
| Estimate Percentage | Archive Job | Determining whether to gather statistics on History tables | Task | BF / Tasks |
| Maximum number of parallel workers | All Jobs | | Job | BF |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|--|--------------------------------|--|----------|--------------------|
| Number of driving table rows per commit for Fully Transactional operation | Archive and Reload Job | NATIVE_MOVE_O2H/H2O - Move data from source tables to history tables and vice versa | Task | BF / Tasks |
| Minimum driving ATT chain size to process in a single commit for Fully Transactional operation | Archive and Reload Job | NATIVE_MOVE_O2H/H2O - Move data from source tables to history tables and vice versa | Task | BF / Tasks |
| Number of rows per commit for Table Parallel copy operation | Archive and Reload Job | NATIVE_MOVE_O2H/H2O - Move data from source tables to history tables and vice versa DELETE_TASK - Delete data from source | Task | BF / Tasks |
| Number of tasks to be created for Data Movement | Archive and Reload Job | PARALLELIZE_MOVE_O2H/H2O - Parallelize moving data from source to history and vice versa | Task | BF / Tasks |
| Selection batch size | Standard and PDM Selection Job | ARCHIVE_SELECTION -Snapshot qualifying OLTP-Table rows into Archive Selection Tables | | |
| Type of Snapshots for Lookup tables | Archive Job | PARALLELIZE_MOVE_O2H/H2O - Parallelize moving data from source to history and vice versa | Task | BF / Tasks |
| Type of indexes to create on History | Archive Job | SEED_CREATE_HISTORY_INDEX - Determining whether to create indexes on History tables | Task | BF / Tasks |
| Type of indexes to drop on History | Archive Job | SEED_CREATE_HISTORY_INDEX - Determining whether to create indexes on History tables | Task | BF / Tasks |

Archive access parameters

The following parameters are located in the archive access section:

| Parameter | Description |
|-------------------------|---|
| Add the rowid column to | Indicates whether the rowid column should be added to the archive |

| Parameter | Description |
|---|--|
| the archive access union view | access union view. |
| Configuration mode | Specifies the configuration mode for archive access. Valid values are UNION and ARCHIVE_ONLY. |
| Create synonyms to dependent packages defined with invokers' rights | Specifies whether dependent packages, which are defined with invokers rights should be cloned, or have synonyms created for them. Creating synonyms can result in a faster runtime for the Create Archive Access job and can result in fewer errors. However, it can also cause problems on some applications. |
| Default optimizer hint for archive access union view | Specifies the default optimizer hint for the archive access union view. |
| Recreate all managed table dependent objects | Indicates whether to recreate objects in reruns of the create archive access schema job. If set to false, reruns will only recreate objects in the archive access schema that need to be recreated. If set to true, then all objects will be recreated. |
| Remote Access Database Link | Specifies the name of the database link in the archive access schema that points to the remote database. |
| Revoke privileges from the archive access schemas | Indicates whether privileges should be revoked from archive access schemas after the archive access run. |

Archive Access parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the Archive access parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|---|----------|--|----------|--------------------|
| Add the rowid column to the archive access union view | CreateAA | CloneObjects0 - Clone Objects, Level 0 | Task | BF / Tasks |
| Configuration mode | CreateAA | CloneObjects0 - Clone Objects, Level 0 | Task | BF / Tasks |
| Create synonyms to | CreateAA | CloneObjects0 - Clone Objects, Level 0 | Task | BF / Tasks |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|--|----------|--|----------|--------------------|
| dependent packages defined with invokers' rights | | | | |
| Default optimizer hint for archive access union view | CreateAA | CloneObjects0 - Clone Objects, Level 0 | Task | BF / Tasks |
| Recreate all managed table dependent objects | CreateAA | SeedDBCloneTask - Initialize | Task | BF / Tasks |
| Remote Access Database Link | CreateAA | CloneObjects0 - Clone Objects, Level 0 | Task | BF / Tasks |
| Revoke privileges from the archive access schemas. | CreateAA | RevokePrivileges | Task | BF / Tasks |

Advanced selection parameters

The following parameters are located in the PDM section:

| Parameter | Description |
|---------------------------|--|
| Materialize Union Views | Indicates whether to materialize union views that are used for data location purposes. For example, during reload, a union view can be used during the selection process to pull data from both the HIST table and the OLTP table. If set to false, union views are not be materialized. If set to true, union views are materialized. |
| Run Output Option | This parameter controls how much PDM produces diagnostic output in the PDM server side log. |
| Selection Parallel Degree | <p>Defines the selection degree of parallelism. Enables parallel processing only if the value is greater than 1. This parallel degree can be overridden at the statement level by putting a hint in <code>sql-tuning.properties</code>. For example:</p> <pre>testenv.ADVANCED_SELECTION...STMT_1.HINT_2=parallel(pc,8) testenv.ADVANCED_SELECTION...STMT_2.HINT_1=parallel(ins,8)</pre> <p>It is not possible to handle parallel degree for each table as the same table can appear in different statements and the same statement can work with different tables, too.</p> |
| Unify MTU Selections | Unify selections in multiple table uses (MTU) into one selection table, and remove duplicate rows. For database to database archiving, the default value of this parameter is true to ensure that rows selected from multiple table uses of the same table are unified and any duplicate rows are removed. If set to false, duplicated rows may be copied into the history table. |

Advanced selection parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the advanced selection parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|---------------------------|-------------------------------|------------------------|---------------|--------------------|
| Materialize Union Views | PDM_* and RELOAD_ PDM jobs | PDM_MERGE_TO_ HIST | Job / Task | |
| Run Output Option | PDM_* and RELOAD_ PDM jobs | Every PDM Task | | |
| Selection Parallel Degree | PDM_* and RELOAD_ PDM jobs | PDM_DTAB_ SELECTION | | BF / Tasks |
| Unify MTU Selections | PDM_* and RELOAD_ PDM jobs | PDM_PROP_ SELECTION | Job / Task | BF / Tasks |

The table below lists the Legend Key and the Legend Value.

Table legend

Legend key: Short name of the task.

Legend value: Description of the task.

| Legend Key | Legend Value |
|--------------------------|-------------------------|
| DP_PARALLEL_DEGREE | DP Parallel Degree |
| MATERIALIZER_UNION_VIEWS | Materialize Union Views |
| RUN_OUTPUT_OPTION | Run Output Option |
| UNIFY_MTU_SELECTIONS | Unify MTU Selections |

Distributed instance parameters

The following parameters are located in the distributed instance section:

| Parameter name | Description |
|---------------------------|---|
| History DB: SID name | The SID name of the history database. |
| History DB: TNS name | The TNS name of the history database. |
| History to OLTP link name | The name of the database link from History to OLTP. |
| OLTP DB: SID name | The SID name of the OLTP database. |
| OLTP DB: TNS name | The TNS name of the OLTP database. |
| OLTP to History link name | The name of the database link from OLTP to History |

Distributed instance parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the distributed instance parameters.

Table legend

Usage: Indicates if it is used by a job, task, or deployment process.

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Usage | Job | Task | Recovery | Pending Jobs/Tasks |
|---------------------------|-------------------------|--|---|----------|--------------------|
| History DB: SID name | Job / Task | Bulk - Interface to Relocation | COPY_DATAFILES - Copy Datafiles from OLTP Server to HIST Server, IMPORT_TTS_RELOC - Run Oracle import-utility to import data from TTS into RELOC Schema | Task | BF / Tasks |
| History DB: TNS name | Job / Task / Deployment | Bulk - Interface to Relocation job, Selection, Archive and Reload jobs | VALIDATE_LA_ENVIRONMENT - Validate Live archive Environment, COPY_DATAFILES - Copy Datafiles from OLTP Server to HIST Server, IMPORT_TTS_RELOC - Run Oracle import-utility to import data from TTS into RELOC Schema | Task | BF / Tasks |
| History to OLTP link name | Deployment | | | | |

| Parameter name | Usage | Job | Task | Recovery | Pending Jobs/Tasks |
|---------------------------|-------------------------|--|--|----------|--------------------|
| OLTP DB: SID name | Job / Task | Bulk - Interface to Relocation | COPY_DATAFILES - Copy Datafiles from OLTP Server to HIST Server, EXPORT_TTS_INTF - Run Oracle export-utility to TTS from INTF Schema | Task | BF / Tasks |
| OLTP DB: TNS name | Job / Task / Deployment | Bulk - Interface to Relocation job, Selection, Archive and Reload jobs | VALIDATE_LA_ENVIRONMENT - Validate Live archive Environment, COPY_DATAFILES - Copy Datafiles from OLTP Server to HIST Server, EXPORT_TTS_INTF - Run Oracle export-utility to TTS from INTF Schema | Task | BF / Tasks |
| OLTP to History link name | Deployment | | | | |

Partition based delete

The following parameters are located in the partition based delete section:

| Parameter name | Description |
|-------------------------------|---|
| Enable Foreign Key Management | Set to true if you want to disable all foreign keys during data movement. Set to false if you do not want to disable the foreign keys during data movement. |
| Partition Based Delete Method | Indicates the method to use for Partition Based Delete, Exchange Partition with an empty table or Truncate partition . |

File naming parameters

The following file naming parameters are used only with database to file archiving and can be set for each individual cartridge:

| Parameter | Description |
|------------------------------|---|
| Group File Prefix | The prefix for the group XML or CSV files. The default value is <cartridge_name>_, where cartridge_name is the name of the cartridge. |
| Group File Suffix | The suffix you want to use for the group XML or CSV files. |
| Group XML Schema File Prefix | The prefix for the group XML schema files. The default value is <cartridge_name>_, where cartridge_name is the name of the |

| Parameter | Description |
|--------------------------------|--|
| | cartridge. |
| Group XML Schema File Suffix | The suffix you want to use for the group XML schema files. |
| Schema Mapping File for Upload | The path and name of a schema mapping file you have created to define the mapping to use when uploading from a file to a destination database. See Map schemas by cartridge, on page 113 . |
| Summary File Prefix | The prefix for the summary files. The default value is <cartridge_name>_, where cartridge_name is the name of the cartridge. |
| Summary File Suffix | The suffix you want to use for the summary files. |
| Summary XML Schema File Prefix | The prefix for the summary XML schema files. The default value is <cartridge_name>_, where cartridge_name is the name of the cartridge. |
| Summary XML Schema File Suffix | The suffix you want to use for the summary XML schema files. |

File naming parameter usage

Each parameter may be used differently by the system. For example, different parameters may be used in different jobs and tasks. In addition, some parameters are used in recovery operations while others are not. The table below lists the usage of the file naming parameters.

Table legend

Job: Name of the job it applies to.

Task: Name of the task it applies to.

Recovery: Indicates if it is picked up by a recovery job or a recovery task.

Pending Jobs/Tasks: Indicates if it is picked up by a pending job or a pending task.

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|-------------------------------------|---------------------|--|----------|--------------------|
| Pending Jobs/TasksGroup File Prefix | COPY_DB_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD. On WebConsole: DB_TO_BE_DIRECT - Extract data from database to backend directly. | Job | BF |
| Group File Suffix | COPY_DB_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD. On | Job | BF |

| Parameter name | Job | Task | Recovery | Pending Jobs/Tasks |
|--------------------------------|---------------------|---|----------|--------------------|
| | | WebConsole: DB_TO_BE_DIRECT - Extract data from database to backend directly | | |
| Group XML Schema File Prefix | COPY_DB_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD On WebConsole: WRITE_ATT_GROUP_XSD - Write XSD file for ATT instance group files | Job | BF |
| Group XML Schema File Suffix | COPY_DB_TO_BE (D2F) | COPY_DB_TO_CSV, CSV2XML, WRITE_ATT_GROUP_XSD On WebConsole: WRITE_ATT_GROUP_XSD - Write XSD file for ATT instance group files | Job | BF |
| Summary File Prefix | COPY_DB_TO_BE (D2F) | WRITE_SUMMARY - Write summary XML and XSD files for XML archive | Job | BF |
| Summary File Suffix | COPY_DB_TO_BE (D2F) | WRITE_SUMMARY - Write summary XML and XSD files for XML archive | Job | BF |
| Summary XML Schema File Prefix | COPY_DB_TO_BE (D2F) | WRITE_SUMMARY - Write summary XML and XSD files for XML archive | Job | BF |
| Summary XML Schema File Suffix | COPY_DB_TO_BE (D2F) | WRITE_SUMMARY - Write summary XML and XSD files for XML archive | Job | BF |

NOTE: All these tasks pickup the new value of the parameter, unless they already have the old value. There is no scenario where you may want to change the value of the parameter during recovery.

Appendix B: Database archiving advanced tasks

This appendix covers advanced tasks that are less frequently performed.

- [SQL tuning, below](#)
- [Gather eligibility analytics for advanced selection, on page 176](#)
- [Reclaim space in the active database, on page 177](#)
- [Use views for customized reporting, on page 178](#)
- [Tune the archive access layer, on page 178](#)
- [Implement partition-based deletion for Oracle, on page 179](#)
- [Enable SQL trace for Oracle, on page 180](#)
- [Add minimal privilege users for Oracle, on page 181](#)
- [Lock and unlock the repository, on page 182](#)

SQL tuning

HPE Structured Data Manager uses SQL hints, session variables, and custom pre-statement execution code to allow you to configure the SQL statements in your Oracle, SQL Server, or Sybase installation, and improve the performance of your selection and data movement statements.

About SQL tuning

You can create hints, set session variables, or create custom pre-statement execution code that affect the selection and data movement statements for your installation.

Each modification is applied to a specific SQL statement for an individual cartridge in a business flow. For example, if the ORDERS_BF business flow contains two cartridges, ORDERS_DOM and ORDERS_INTL, you would have to write separate hints for each cartridge, or use wildcards.

Each cartridge contains the following seed statement files in the cartridge directory:

| Seed statement files | Description |
|-------------------------------|--|
| seed-selection-statements.xml | Contains SQL statements that apply to the cartridge selection process. Data selection statements are supported for database to database cartridges and database to file cartridges. |
| seed-dm-sql-statements.xml | Contains SQL statements that apply to the cartridge data movement process. Data movement statements are supported for database to database |

| Seed statement files | Description |
|----------------------|-------------|
| | cartridges. |

With wildcards, each modification can affect one or more statements in each seed statement file. The modifications are defined in the `sql_tuning.properties` file, which is stored in the home directory under `config`. See [Manage the home directory, on page 39](#).

If you are using Advanced Selection, selection statements are created at runtime. If you want to use SQL tuning with Advanced Selection, you should run Advance Selection in diagnostic mode by setting the configuration parameter Run Output Option to `SHOW_PLAN` to get a list of all selection statements that Advanced Selection will execute along with the associated execution plans for each statement. You can use the diagnostic information to create specific SQL tuning hints.

About statements

When you deploy database to file or database to database archiving, the SQL statements in the seed-selection-statements.xml and seed-dm-sql-statements.xml files are generated. Each statement contains the following information:

- source environment
- context
- appspack_name
- table_identifier
- statement_name

Example

```
<SRC_ENV>Oracle_OLTP</SRC_ENV>
<CONTEXT>OLTP_SELECTION</CONTEXT>
<APPSPACK_NAME>d2d_trans</APPSPACK_NAME>
<TABLE_IDENTIFIER>ORDER_HEADER</TABLE_IDENTIFIER>
<STATEMENT_NAME>INSERT_SELECTION_NO_ANALYTICS
</STATEMENT_NAME>
```

About context values

The following context values are valid:

| Context | Description |
|--------------------|--|
| OLTP_SELECTION | The archive selection step. |
| ADVANCED_SELECTION | The advanced selection step. (Use this context instead of OLTP_SELECTION when using Advanced Selection.) |

| Context | Description |
|----------------------------|---|
| ELIGIBILITY_ ANALYTICS | Eligibility analytics step. |
| NATIVE_MOVE_ O2H | Archive copy step. Data is copied from the source (OLTP) database to the target (history) database, and then deleted. For database to database archiving only. |
| HISTORY_ SELECTION | History selection step for reload. For database to database archiving only. |
| NATIVE_MOVE_ H2O | History copy step for reload. Data is copied from the target (history) database to the source (OLTP) database and then deleted. For database to database archiving only. |
| ARCHIVE_UNDO_ SELECTION | Selection step for archive undo. For database to database archiving only. |
| NATIVE_COPY_ O2I | Copy OLTP to interface step. Data is copied from the OLTP to the interface tables. For database to database archiving bulk archive only. |
| NATIVE_COPY_ R2H | Copy relocation to history step. Data is copied from the relocation tables to the history tables. For database to database archiving bulk archive only. |
| NATIVE_COPY_ O2RT | Copy OTLP to retention step. For database to database archiving bulk archive only. |
| NATIVE_COPY_ RT2O | Copy retention to OLTP step. For database to database archiving bulk archive only. |

About statement_name values

Each context has one or more of the following statement names.

Selection statement names

The following statement values are used when inserting into the selection tables:

| Statement | Description |
|-------------------------------|--|
| INSERT_SELECTION_ANALYTICS | Used when the Eligibility Analytics Configuration configuration parameter in the Web Console is set to “Enable the eligibility analytics”. Eligibility Analytics must be enabled in both Designer and the Web Console before eligibility analytics occurs. |
| INSERT_SELECTION_NO_ANALYTICS | Used when the Eligibility Analytics Configuration configuration parameter in the Web Console is set to “Disable the eligibility analytics”. |
| STMT_x | The format of statement names for Advanced Selection where x is an integer relating to the target statement number. You can obtain the list of statements by running Advanced Selection in diagnostic mode and setting the configuration parameter RUN_OUTPUT_OPTION to SHOW_PLAN. |

Copy and move statement names

The following statement values are used when moving data between the source database and the target. Each move operation consists of an INSERT statement and a DELETE statement.

| Statement | Description |
|---------------------|---|
| COPY_TRX_RANGE | Used for fully transactional data moment when RANGE is applicable. |
| COPY_TRX_NO_RANGE | Used for fully transactional data moment when NO_RANGE is applicable. |
| COPY_TP_RANGE | Used for table parallel data moment when RANGE is applicable. |
| COPY_TP_NO_RANGE | Used for table parallel data moment when NO_RANGE is applicable. |
| DELETE_TP_RANGE | Used for deletion with table parallel when RANGE is applicable. |
| DELETE_TRX_RANGE | Used for deletion with fully transactional when RANGE is applicable. |
| DELETE_TP_NO_RANGE | Used for deletion with table parallel when NO_RANGE is applicable. |
| DELETE_TRX_NO_RANGE | Used for deletion with fully transactional when NO_RANGE is applicable. |

NO_RANGE will be executed when it is a single task per driving table (fully transactional) or table (table parallel), and one of the following configuration combinations is set:

| Configuration parameter | Required value |
|---|----------------|
| Use database parallelism to archive or reload data | TABLE_PARALLEL |
| Number of rows per commit for Table Parallel copy operation | 0 |

or

| Configuration parameter | Required value |
|---|---------------------|
| Use database parallelism to archive or reload data | FULLY_TRANSACTIONAL |
| Number of driving table rows per commit for Fully Transactional operation | 0 |

If neither of these two combinations is set, RANGE will be used.

Eligibility analytics statement name

The following statement value is used for eligibility analytics:

| Statement | Description |
|-----------------------|---------------------------------|
| ELIGIBILITY_ANALYTICS | Used for eligibility analytics. |

Archive undo statement names

The following statement values are used for the archive undo job:

| Statement | Description |
|-----------------------------------|--|
| CHECK_SELECTION_FOR_ARCHIVE_UNDO | Used for checking if the archive undo selection process is required for a specific table. If the archive undo job failed the last time it was run, some of the tables may have been successfully processed. This statement skips any tables that do not need to be added to the selection tables. |
| INSERT_SELECTION_FOR_ARCHIVE_UNDO | Used for loading selection tables with information about the rows that need to be processed by the archive undo job. |

About appspack_name and table_identifier values

The appspack_name and table_identifier values depend on your cartridge and the tables included in that cartridge. The appspack_name corresponds to the cartridge name designated in Designer, and the table_identifier corresponds to the table alias designated in the model in Designer.

See the seed-selection-statements.xml and seed-dm-sql-statements.xml files for the valid values.

In Advanced Selection, the table_identifier values must be set using the wildcard operation because specific table name hints are not supported in advanced selection.

About using session variable statements

Session variable statements require different commands depending on the database you are using.

| Database | Command |
|------------|---|
| Oracle | Uses the ALTER SESSION command arguments. |
| SQL Server | Uses the SET command arguments. |

Because HPE Structured Data Manager automatically wraps it with the appropriate command (ALTER SESSION or SET), the variable value is just the arguments for the database command.

SQL Server example

For SQL Server, if you want to set the deadlock priority to low, and you would normally change the settings with:

```
SET deadlock_priority low
```

set the value to:

```
deadlock_priority_low
```

Oracle example

For Oracle, if you want to set optimizer index caching to 10, and you would normally change the optimizer settings with:

```
ALTER SESSION set optimizer_index_caching=10
```

set the value to:

```
set optimizer_index_caching=10
```

See [Edit the sql_tuning.properties file, on the next page](#)

About using custom pre-statement execution code

The PreExec code enables you to specify custom code to be run before the selection or data movement statements.

Pre-statement execution code is run in different locations depending on the type of statement.

| Statement | Location |
|---------------------|--|
| Selection statement | The PreExec code is run on the source local database. |
| Copy statement | The PreExec code is run on the target local database. |
| Delete statement | For Oracle, the PreExec code is run on the remote source database. For other databases, the code is run on the remote target database. |

NOTE: If the statement you are modifying is for a Reload job, the locations are reversed.

Example

```
Oracle_OLTP.OLTP_SELECTION.INVENTORY.ITEM.INSERT_SELECTION_ANALYTICS.PREEXEC_
NAME=apps_initialization
```

where `apps_initialization` is the name of a SQL block that you want to execute.

For advanced selection, the syntax for the session variable is:

```
<src_env>.ADVANCED_SELECTION.<appspack_name>.*.*.SESSION_VAR1=set optimizer_index_
caching = 1
```

where `TABLE_IDENTIFIER` and `STATEMENT_NAME` are set to wildcards.

Edit the `sql_tuning.properties` file

Use the `sql_tuning.properties` file to add SQL hints or set session variables.

1. Navigate to the directory under the home directory that contains the seed statements for the cartridge you want to modify. See [Manage the home directory, on page 39](#).

For example:

```
/artifacts/businessflow/<busflow_name>/cartridge/<product>/<cartridge_name>
```

| Parameter | Description |
|----------------------|--|
| <install_directory> | The location where you installed the software |
| <business_flow_name> | The name of the business flow that contains the cartridge you want to modify |
| <product> | The archiving method the cartridge uses: <ul style="list-style-type: none">• oa—database to file archiving• la—database to database archiving |
| <cartridge_name> | The individual cartridge you want to query |

NOTE: Database to file cartridges only support selection statement hints.

2. Open the seed statement file you want to view.
 - `seed-selection-statements.xml`—contains selection statements
 - `seed-dm-sql-statements.xml`—contains data movement statements

NOTE: Do not edit the seed statement files.

3. Search the seed statement file for the statement you want to modify.

You need to note the following information:

- source environment
- context
- `appspack_name`
- `table_identifier`
- `statement_name`

Example

```
<SRC_ENV>Oracle_OLTP</SRC_ENV>  
<CONTEXT>OLTP_SELECTION</CONTEXT>  
<APPSPACK_NAME>d2d_trans</APPSPACK_NAME>  
<TABLE_IDENTIFIER>ORDER_HEADER</TABLE_IDENTIFIER>  
<STATEMENT_NAME>INSERT_SELECTION_NO_ANALYTICS  
</STATEMENT_NAME>
```

- 4. In a text editor, open the sql_tuning.properties file, which is stored in the home directory under config. See [Manage the home directory, on page 39](#).

NOTE: For Oracle, the sql_tuning.properties file has default hints. The SQL Server sql_tuning.properties file has no defaults.

To add a hint

- 5. Add a new hint using the following format:

```
<SRC_ENV>.<CONTEXT>.<APPSPACK_NAME>.<TABLE_IDENTIFIER>.<STATEMENT_<br>NAME>.<hintType>=<hint>
```

| Parameter | Description |
|--------------------|---|
| <SRC_ENV> | The contents of the <SRC_ENV> tag in the seed statement file. |
| <CONTEXT> | The contents of the <CONTEXT> tag in the seed statement file. |
| <APPSPACK_NAME> | The contents of the <APPSPACK_NAME> tag in the seed statement file. |
| <TABLE_IDENTIFIER> | The contents of the <TABLE_IDENTIFIER> tag in the seed statement file. |
| <STATEMENT_NAME> | The contents of the <STATEMENT_NAME> tag in the seed statement file. |
| <hintType> | INSERT_HINT, SELECT_HINT or DELETE_HINT. |
| <hint> | <p>The SQL hint.</p> <p>In Advanced Selection, the hint name is defined using the following format:</p> <p>HINT_x</p> <p>where x is an integer. You can obtain the hint number from the list of statements produced by running Advanced Selection in diagnostic mode, setting the configuration parameter "RUN_OUTPUT_OPTION" to "SHOW_PLAN".</p> |

Example hint

```
OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_  
ANALYTICS.<hintType>=<hint>
```

TIP: You can use wildcards to apply the changes to more than one statement.
..*.*.<hintType>=<hint> applies the hint to all statements in all installed cartridges.

Example SQL in seed file

A SQL statement in a seed file looks something like the following sample. The items between the ## symbols are the available hintTypes for that statement.

```
INSERT ##INSERT_HINT## INTO "CUSTOLTTP_HIST"."ORD" ("ORDERDATE", "ORDID", "OBT_  
WF_RUN_ID", "OBT_ROW_SEQ", "OBT_TIMESTAMP", "OBT_SAVED_ROWID" ) SELECT  
##SELECT_HINT## "S"."ORDERDATE",  
    "S"."ORDID",  
    "S"."OBT_WF_RUN_ID",  
    "S"."OBT_ROW_SEQ",  
    sysdate,  
    "S"."OBT_SAVED_ROWID"  
FROM "OBT_IF"."ORD_SVO"@ORCL3 "S"  
WHERE ("S"."OBT_WF_RUN_ID"=?)  
AND ("S"."OBT_ROW_SEQ" BETWEEN ? AND ?)
```

6. Fill in the appropriate SQL hint.

Oracle example

```
Oracle_OLTP.OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_  
ANALYTICS.SELECT_HINT=use_n1
```

SQL Server example

```
MSOLTP.OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_  
ANALYTICS.SELECT_HINT=FAST 5
```

7. Save the sql_tuning.properties file.

To add a session variable

8. Add a new session variable using the following format:

```
<SRC_ENV>.<CONTEXT>.<APPSPACK_NAME>.<TABLE_IDENTIFIER>.<STATEMENT_  
NAME>.<SESSION_VARn>=<var>
```

| Parameter | Description |
|-----------------|---|
| <SRC_ENV> | The contents of the <SRC_ENV> tag in the seed statement file. |
| <CONTEXT> | The contents of the <CONTEXT> tag in the seed statement file. |
| <APPSPACK_NAME> | The contents of the <APPSPACK_NAME> tag in the seed statement file. |

| Parameter | Description |
|--------------------|--|
| <TABLE_IDENTIFIER> | The contents of the <TABLE_IDENTIFIER> tag in the seed statement file. |
| <STATEMENT_NAME> | The contents of the <STATEMENT_NAME> tag in the seed statement file. |
| <SESSION_VARn> | The identifier for the session variable. |
| <var> | The session variable. |

Example

Oracle_OLTP.OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_ ANALYTICS.<SESSION_VARn>=<var>

TIP: You can use wildcards to apply the changes to more than one statement. *. *. *. *.<SESSION_VARn>=<var> applies the session variable to all statements in all installed cartridges.
For example, if the default statement is as follows: DefEnv.OLTP_SELECTION. *. *. *.INSERT_HINT=APPEND
changing the statement to: DefEnv.OLTP_SELECTION.INVENTORY.ITEM. *.INSERT_HINT=PARALLEL
will override the APPEND hint for only the ITEM table in INVENTORY.

- Fill in the appropriate session variable.

Oracle example

Oracle_OLTP.OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_ ANALYTICS.SESSION_VAR3=set optimizer_index_caching=10

SQL Server example

MSOLTP.OLTP_SELECTION.d2d_trans.ORDER_HEADER.INSERT_SELECTION_NO_ ANALYTICS.SESSION_VAR3=deadlock_priority low

- Save the sql_tuning.properties file.

To add a preexec statement

- Add a new pre-statement execution code using the following format:

<SRC_ENV>.<CONTEXT>.<APPSPACK_NAME>.<TABLE_IDENTIFIER>.<STATEMENT_NAME>.<PREEXECNAME>=<preexecname>

| Parameter | Description |
|-----------------|---|
| <SRC_ENV> | The contents of the <SRC_ENV> tag in the seed statement file. |
| <CONTEXT> | The contents of the <CONTEXT> tag in the seed statement file. |
| <APPSPACK_NAME> | The contents of the <APPSPACK_NAME> tag in the seed statement file. |

| Parameter | Description |
|--------------------|--|
| <TABLE_IDENTIFIER> | The contents of the <TABLE_IDENTIFIER> tag in the seed statement file. |
| <STATEMENT_NAME> | The contents of the <STATEMENT_NAME> tag in the seed statement file. |
| <preexecname> | The name of the custom code you want to execute. |

Example

```
Oracle_OLTP.OLTP_SELECTION.INVENTORY.ITEM.INSERT_SELECTION_ANALYTICS.PREEXEC_
NAME=apps_initialization
```

where apps_initialization is the name of the custom code you want to execute.

TIP: You can use wildcards to apply the changes to more than one statement.
*. *.*.*.*.<PREEXEC_NAME>=<preexecname> applies the modification to all statements in all installed cartridges.

12. Save the sql_tuning.properties file.

Specify SQL hints for non-intrusive environments

In the sql_tuning.properties file, you can add hints for optimizing particular SQL statements. You can specify hints for driving statements and cursor loops.

Hints for driving statements have the following syntax:

```
<env_name>.OLTP_LITE_SELECTION.<cartridge>.<table_alias>.
DRIVING_STATEMENT.<hint_type>=<hint>
```

Hints for cursor loops have the following syntax:

```
<env_name>.OLTP_LITE_SELECTION.<cartridge>.<table_alias>.
CURSOR_LOOP.<hint_type>=<hint>
```

NOTE: You should understand how your particular database handles hints before implement them in your sql_tuning.properties file.

| Parameter | Description |
|---------------|--|
| <env_name> | The database that the hints is being applied to. For example, SQL Server or Oracle. |
| <cartridge> | The name of the cartridge. You can specify * to apply the hint to all cartridges. |
| <table_alias> | The table to apply the hint to. You can specify * to apply the hint to all tables. |
| <hint_type> | The hint type specifies where in the SQL statement to insert the hint. The possible values are: <ul style="list-style-type: none">• PREFIX-HINT—The hint is at the beginning the SQL statement before the SELECT |

| Parameter | Description |
|-----------|---|
| | <p>clause.</p> <ul style="list-style-type: none"> • SELECT_HINT—The hint is placed after SELECT in the SQL statement. • TABLE_ALIAS_HINT—The hint is placed after each table alias in the FROM clause. • SUFFIX_HINT—The hint is placed at the end of the SQL statement. |
| <hint> | <p>In Advanced Selection, the hint name is defined using the following format:</p> <p>HINT_x</p> <p>where x is an integer. You can obtain the hint number from the list of statements produced by running Advanced Selection in diagnostic mode, setting the configuration parameter "RUN_OUTPUT_OPTION" to "SHOW_PLAN".</p> |

Specify SQL hints for advanced selection

In the `sql_tuning.properties` file, if you have not customized the default hints for advanced selection, you should comment them out with the `#` character as shown. You need to make this change because the previous default might disallow the Web Console configuration parameter for parallel degree for statements 9, 10, and 11.

NOTE: Advanced selection does not support the use of wildcards for statement names and hint types. Specific statement names and specific hint types should be defined in order for the hint to be used.

```
#=====
# Default Advanced Selection Hints for Oracle
# Note: Advanced selection does not support the use of
# wildcards for statement names and hint types.
# Specific statement names and specific hint types should be
# defined in order for the hint to be used.
#DefEnv.ADVANCED_SELECTION...STMT_9.HINT_2=parallel(fr_s)
# parallel(c) parallel(pia)
#DefEnv.ADVANCED_SELECTION...STMT_10.HINT_1=parallel(fr_s)
# parallel(src) parallel(tgt) parallel(to_s)
#DefEnv.ADVANCED_SELECTION...STMT_11.HINT_2=parallel(c1)
# parallel(npia) parallel(pia)
```

The parallel hint for advanced selection needs the table alias and the degree, unlike above default statements:

```
#DefEnv.ADVANCED_SELECTION...STMT_11.HINT_2=parallel(c1,8)
# parallel(npia,8) parallel(pia,8)
```

TIP: You need not redeploy for the hints to take effect.

Gather eligibility analytics for advanced selection

When performing eligibility analytics for standard selection, you can follow the steps in [View eligibility analytics, on page 58](#). For advanced selection, you must perform some additional queries against the advanced selection tables to determine which rows are ineligible for archiving.

To ascertain ineligible rows for advanced selection

1. Find the controller id:

```
select * from pdm_controllers where controller_name =<pdm_cartridge_name>
```

2. Find the alias for the driving table:

```
select * from pdm_nodes where controller_id= <controller_id>
```

3. Get the node and cartridge exception table:

```
select 'node' || node_id || '_s_' || c.controller_id NODE, 'PDM_CARTRIDGE_EXC_' ||  
c.controller_id EXCEPTION  
from pdm_controllers c,  
pdm_nodes n  
where c.controller_id = n.controller_id  
and c.controller_name =<pdm_cartridge_name>  
and n.alias = <alias_name>
```

Assume for the purposes of the example that the driving table for your model is OE_ORDER_HEADERS_ALL.

4. Obtain the total driving table rows:

```
select count(*) from pdm_chains
```

5. Obtain the total ineligible driving table rows:

```
select count(*) from pdm_chains where at_id in (select at_id from pdm_invalid_ats)
```

6. Find the at_id, if you have a specific header_id, for example, the primary key for OE_ORDER_HEADERS_ALL:

```
select at_id from node<n>_s_<controller_id> a  
where a.header_id = &header_id  
and exists (select 'x' from pdm_invalid_ats b  
where a.at_id = b.at_id)
```

where node<n>_s_<controller_id> is replaced with the result from [Get the node and cartridge exception table](#); above.

7. Using &at_id from the results of the previous step and PDM_CARTRIDGE_EXC_<controller_id> from the results of [Get the node and cartridge exception table](#); above, find out why certain driving table rows are ineligible:

```
select * from PDM_CARTRIDGE_EXC_<controller_id>  
where at_id = &at_id
```

8. Advanced selection populates a table named pdm_table_exc_<controller_id>_<table_id>. Using controller_id from [Find the controller id](#); above and table_id from [Using &at_id from the results of the previous step and PDM_CARTRIDGE_EXC_<controller_id>](#) from the results of [Get the node and cartridge exception table](#); above, find out why certain driving table rows are

ineligible; above, you can query this table to get the actual rows:

- a. Run this query to retrieve obt_saved_rowid

```
select * from pdm_table_exc_<controller_id>_<table_id>
where row_id in (select row_id
from PDM_CARTRIDGE_EXC_<controller_id>
where at_id = &at_id)
```

where &at_id comes from the results of [Find the at_id, if you have a specific header_id, for example, the primary key for OE_ORDER_HEADERS_ALL;](#), on the previous page.

- b. You can now further join to table_name from [Find the at_id, if you have a specific header_id, for example, the primary key for OE_ORDER_HEADERS_ALL;](#), on the previous page to get the rows:

```
select b.*
from pdm_table_exc_<controller_id>_<table_id> a, <table_name> b
where row_id in (select row_id from PDM_CARTRIDGE_EXC_<controller_id>
where at_id = &at_id) and a.obt_saved_rowid = b.rowid;
```

where &at_id comes from the results of [Find the at_id, if you have a specific header_id, for example, the primary key for OE_ORDER_HEADERS_ALL;](#), on the previous page.

Reclaim space in the active database

The amount of disk space necessary for the active database is reduced after the inactive archived data has been segregated and relocated.

After running a business flow, you should perform the following types of tasks using your company's procedures to reclaim space in the active database.

To reclaim space after running the Transactional Archive

Oracle

1. Perform a database reorganization.
2. Perform a datafile reorganization.

SQL Server

1. Use the DBCC SHRINKDATABASE command or DBCC CLEANTABLE command to reclaim space in the database.

For more information, see the Oracle or SQL Server documentation.

To reclaim space after running Bulk Archive

1. Perform a datafile reorganization.

The tables in the active database are automatically reorganized as part of the Bulk Archive process.

For more information, see your Oracle documentation.

Use views for customized reporting

HPE Structured Data Manager supports the following views of the product metadata:

| View | Definition |
|--------------------------|--|
| OBTWC_BF_RUNS_V | Displays business flow runs, but does not display the detailed steps. |
| OBTWC_RUN_DETAILS_V | Displays business flow runs and the detailed steps of the business flow. |
| OBTWC_RUN_PARAMETERS_V | Displays the name and value of each parameter and maps them to an individual job run_id. |
| OBTWC_ROWCOUNT_DETAILS_V | Display the row count and row count operation for each table. |

You can use these views to create customized reports for your installation.

Tune the archive access layer

To improve the performance of the archive access layer for data transparency (database to database), HPE Structured Data Manager provides a number of parameters. By adding a logon trigger to the archive access schema, you can set these parameters to tune performance in your environment.

NOTE: You need to add the logon trigger to any archive access schema used to access the combined data. If you only have one archive access schema, you would add it just to that schema. If you used multiple archive access schemas to combine the data, then you would need to add it to each of those schemas.

TIP: The Oracle E-Business Suite integration kit automatically installs a logon trigger in the archive access schema for you. For PeopleSoft integration kits, you need create the trigger yourself.

Example

You could add a logon trigger similar to the following to all of the archive access schemas from which you are combining data:

```
CREATE or replace TRIGGER obtva_cust_logon AFTER logon ON apps_aa.schema
begin
declare
    procedure execute_sql (p_sql_stmt in varchar2) is
        l_count      number;
        l_c          number;
        l_rc         number;
    begin
        l_c := dbms_sql.open_cursor;
        dbms_sql.parse(l_c, p_sql_stmt, dbms_sql.native);
```

```
        l_rc := dbms_sql.execute(l_c);
        dbms_sql.close_cursor(l_c);
    end execute_sql;
BEGIN
execute_sql('ALTER SESSION SET "_PUSH_JOIN_PREDICATE"=TRUE');
execute_sql('ALTER SESSION SET "_COMPLEX_VIEW_MERGING"=true');
execute_sql('ALTER SESSION SET "_UNNEST_SUBQUERIES"=FALSE');
execute_sql('ALTER SESSION SET optimizer_index_cost_adj=100');
execute_sql('ALTER SESSION SET optimizer_index_caching=0');
execute_sql('ALTER SESSION SET "_index_join_enabled"=false');
execute_sql('ALTER SESSION SET "_optimizer_push_pred_cost_based"=false');
execute_sql('ALTER SESSION SET "_optimizer_cost_based_transformation"=on');
execute_sql('ALTER SESSION SET "_push_join_union_view"=true');
end;
END;
/
```

Implement partition-based deletion for Oracle

When you create your cartridge in the Designer, you can choose to use partition-based deletion, if the source is Oracle 11c or 12g, and all rows are selected.

- Partition-based deletion has some system parameters that you can set. See [Partition based delete , on page 161](#).

NOTE: These values must be set prior to running the job.

- The system creates the exchange table, performs the partition exchange, then drops the exchange table. However, if you create the exchange table yourself manually, then you must register it using the API `addXChangeTableMap.groovy`.

Once the table is registered it will not be dropped after the partition exchange, but it will be truncated.

To un-register the table, use the API `removeXChangeTableMap.groovy`.

- Oracle provides several options for partition exchange statements, like `INCLUDING INDEXES` and `WITHOUT VALIDATIONS`. You can apply such options in `sql_tuning.properties`. See [SQL tuning, on page 164](#) for more information about `sql_tuning.properties`.

For example, adding the line below, controls the options that should apply during partition exchange, and hence the system will execute:

```
DefaultEnvironment.PBD_H2O.INVENTORY.ORDER.PARTITION_EXCHANGE.STMT_EXTN=INCLUDING
INDEXES WITHOUT VALIDATION UPDATE GLOBAL INDEXES
ALTER TABLE <PARTITIONED_TABLE> EXCHANGE PARTITION WITH TABLE <XCHANGETABLE>
INCLUDING INDEXES WITHOUT VALIDATION UPDATE GLOBAL INDEXES
```

- During the deployment, the history table is created with the same partitioning as the corresponding source table. However, you can recreate the history table and indexes manually, if any extra storage attribute needs customization after the deployment or, before the deployment, when the history schema mapping is set using the API `addHistoryOwnerMapping.groovy` prior to deployment.
- If rows are selected from more than one partition, a partition-based delete is performed on all the qualifying partitions.

- After the partition exchange, any unusable global indexes are rebuilt and all local indexes are rebuilt on the empty partition to ensure the indexes are in valid state.

The only valid data movement for partition-based delete is Table Parallelism.

If you set the **Data Movement Method** parameter to **Move Data in a Fully Transactional manner**, that setting is ignored by the job. For Reload, the run uses the fully transactional method for all tables except those marked for partition-based delete, where the management of foreign keys comes into the play.

Enable SQL trace for Oracle

You can enable or disable Oracle SQL trace using properties files. Because SQL trace can impact performance, HPE recommends that you only enable it for performance tuning.

- To enable SQL tracing for the database to database archiving and database to file selection, use the `sql_tuning.properties` file.
- To enable SQL tracing for database to file data movement, use the `outerbay.properties` file.

The trace files are saved to the `user_dump_dest` directory. For more information, see your Oracle documentation.

Enable SQL trace using the `sql_tuning.properties` file

Use the `sql_tuning.properties` file to enable SQL trace for database to database selection and data movement, and database to file selection.

1. Navigate to `config` under the home directory. See [Manage the home directory, on page 39](#).
2. Open the `sql_tuning.properties` file. See [SQL tuning, on page 164](#).
3. Add a session variable to enable SQL trace.

Example

```
SESSION_VAR1=set events '10053 trace name context forever, level 1'
```

The trace begins before the statement that the session variable is attached to.

For example, if you attach it to the following statement:

```
OLTP_SELECTION.INVENTORY.ITEM.INSERT_SELECTION_ANALYTICS.SESSION_VAR1=set  
events '10053 trace name context forever, level 1'
```

the trace command is issued before the `OLTP_SELECTION.INVENTORY.ITEM. INSERT_SELECTION_ANALYTICS` SQL statement is run.

Enabling SQL trace using the `outerbay.properties` file

Use the `outerbay.properties` file to enable SQL trace for database to file data movement.

4. Navigate to the `outerbay.properties` file, which is stored in the home directory under `config`. See [Manage the home directory, on page 39](#).
5. Open the `outerbay.properties` file with a text editor.
6. Locate the SQL trace section.
7. Update the properties file. By default, the following properties are set:

```
DASL.sqlTrace.level=1  
DASL.sqlTrace.components=ALL
```

8. Save the outerbay.properties file.

Add minimal privilege users for Oracle

An ARCHADMIN user is a user that has the minimum set of privileges for performing archive tasks. You create an ARCHADMIN user when you want to avoid having an archive administrator that has a database administrator role, making sure that data security is not compromised. The ARCHADMIN user can only deploy business flows. This user cannot be used for:

- Installing repositories
- Installing or uninstalling environments
- JDBC connections
- Databases other than Oracle
- Packaged applications (such as PeopleSoft)

The ARCHADMIN user is created as either a source ARCHADMIN user or target ARCHADMIN user. Source ARCHADMIN users are created in the source database (OLTP). Target ARCHADMIN users are created in the target database (History). The target database ARCHADMIN user is required for distributed instance configurations.

The script or scripts that you run to create an ARCHADMIN user depend on the configuration for cartridge deployment. The scripts support the following configurations:

- Database to file
- Standard selection
- Advanced selection
- Database to database
- Standard SI and DI
- Advanced selection DI
- Advanced PDM SI and DI
- Advanced bulk archive

NOTE: An ARCHADMIN user does not support deployment of cartridges for database to file in non-intrusive environments.

To create an ARCHADMIN user

1. Open a command window by selecting **Start>Run**, typing `cmd`, then clicking **OK**.
2. Change to the directory to `<install_directory>/obt/foundation/util/sql/oracle`.
3. Connect to SQL*Plus with the database administrator user having the DBA role.
4. From the SQL*Plus command line, run the appropriate script or scripts based on your configuration for cartridge deployment.

| Standard Selection Configurations | Script |
|-------------------------------------|--|
| Database to file archive | obt_source_admin_user_d2f_bf.sql |
| Transactional data movement archive | obt_source_admin_user_d2d_bf.sql obt_target_admin_user_d2d_bf.sql |
| Transactional data movement reloads | obt_source_admin_user_d2d_bf.sql obt_target_admin_user_d2d_bf.sql |

| Advanced Selection Configurations | Script |
|-------------------------------------|--|
| Transactional data movement archive | obt_source_admin_user_d2d_bf.sql obt_target_admin_user_d2d_bf.sql |
| Transactional data movement reload | obt_source_admin_user_d2d_bf.sql obt_target_admin_user_d2d_bf.sql |

| Upload Configuration | Script |
|-------------------------|--------------------------------------|
| File to Database upload | obt_source_admin_user_f2d_upload.sql |

- At the prompts enter the user name for the ARCHADMIN user, source or target administrator user password, and tablespace name.

NOTE: You can have more than one ARCHIVEADMIN user in the source and target databases. However, each user must have a different name. For example, ARCHADMIN_1, ARCHADMIN_2, and so on.

- Exit SQL*Plus.
- Deploy your business flow using the ARCHADMIN user name and password for the source database administrator. For distributed instances, use the ARCHADMIN user name and password for the target database administrator. See the Web Console online help and [Deploy and run business flows, on page 43](#) for more information about deploying business flows in the Web Console.

Lock and unlock the repository

In some situations, you may need to lock the repository. When you lock the repository, you can:

- prevent business flows and cartridges from being deployed in all environments.
- prevent business flows and cartridges from being deployed in a particular environment.
- prevent any new environments from being created.
- prevent any existing environments from being deleted.

For example, you may want to stop all deployments while you perform some system maintenance or upgrade activities. After you are ready, you can then unlock the repository to allow business flows and cartridges to again be deployed.

NOTE: All new environments are unlocked by default.

Lock the repository

To lock the repository

NOTE: If a business flow or cartridge deployment is already in process, locking the repository has no effect on that deployment.

1. Navigate to the <install_directory>/obt/bin directory.
where <install_directory> is the location where you installed the software.
2. Run the lock_repository script using the appropriate syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <code>./lock_repository.sh [-m <lock_mode>] [-e <environment_name>] -u <yourName> [-c <comment>]</code> |
| DOS | <code>lock_repository.bat [-m <lock_mode>] [-e <environment_name>] -u <yourName> [-c <comment>]</code> |

| Parameter | Description |
|------------------|---|
| lock_mode | Optional. If you want to specify a lock mode, enter NO_CREATE_REPOS. This lock mode prevents any environments from being created or deleted. |
| environment_name | Optional. The name of the environment. |
| yourName | The name of the person locking the repository. Used for auditing purposes only. If name contains spaces, you need to use double-quotes around the name. |
| comment | Optional. The reason why the repository is being locked. |

3. Enter the encryption key at the prompt.

Windows examples

- To prevent business flows and cartridges from being deployed in all environments:
`lock_repository.bat -u Joe`
- To prevent business flows and cartridges from being deployed in the environment Oracle_Env:
`lock_repository.bat -e Oracle_Env -u "Joe Smith"`
- To prevent new environments from being created:
`lock_repository.bat -m NO_CREATE_ENV -u Paul -c "No new environment creation"`

Unlock the repository

To unlock the repository

1. Navigate to the `<install_directory>/obt/bin` directory.
where `<install_directory>` is the location where you installed the software.
2. Run the `unlock_repository` script using the appropriate syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <code>./unlock_repository.sh [-m <lock_mode>] [-e <environment_name>] -u <yourName> [-c <comment>]</code> |
| DOS | <code>unlock_repository.bat [-m <lock_mode>] [-e <environment_name>] -u <yourName> [-c <comment>]</code> |

| Parameter | Description |
|-------------------------------|--|
| <code>lock_mode</code> | Optional. If you want to specify a lock mode, enter <code>NO_CREATE_REPOS</code> . This lock mode enables environment creation and deletion in the repository. |
| <code>environment_name</code> | Optional. The name of the environment. |
| <code>yourName</code> | The name of the person unlocking the repository. Used for auditing purposes only. If name contains spaces, you need to use double-quotes around the name. |
| <code>comment</code> | Optional. The reason why the repository is being unlocked. |

The script prompts for the encryption key.

3. Enter the encryption key at the prompt.

Windows examples

- To unlock all environments:
`unlock_repository.bat -u Paul -c "Unlock to allow check-ins"`
- To unlock just the environment `Oracle_Env`, without affecting any other locked environments:
`unlock_repository.bat -e Oracle_Env -u Mike -c "Unlock Oracle_Env"`
- To allow new environments to be created:
`unlock_repository.bat -m NO_CREATE_ENV -u Paul -c "Allow environment creation"`

Appendix C: Change encryption keys and passwords

The password manager utility is used to change the encryption key and repository password.

- [Change the encryption key, below](#)
- [Reset the encryption key, below](#)
- [Change the repository password, on the next page](#)
- [Update database location passwords, on page 187](#)

Change the encryption key

Use password manager to change the encryption key for the environment.

To change the encryption key

1. Open a command window.
2. Run the utility using the following syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <install_directory>/obt/bin/runPwManager.sh -m |
| DOS | <install_directory>/obt/bin/runPwManager.bat -m |

where <install_directory> is the location where you installed the software.

3. Enter values for the prompts. Enter a question mark (?) to see a list of available values.

- Enter environment to modify:
- Enter CURRENT encryption key:
- Re-enter CURRENT encryption key:
- Enter NEW encryption key:
- Re-enter NEW encryption key:

All encrypted passwords in the connection-sources.xml file are re-encrypted using the new key. The encryption key values are not displayed on the window.

Reset the encryption key

Use password manager to reset a lost encryption key for the environment.

To reset the encryption key

1. Open a command window.
2. Run the utility using the following syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <install_directory>/obt/bin/runPwManager.sh -s |
| DOS | <install_directory>/obt/bin/runPwManager.bat -s |

where <install_directory> is the location where you installed the software.

3. Enter values for the prompts. Enter a question mark (?) to see a list of available values.
 - Enter environment to modify:
 - Enter repository database administrator password:
 - Enter NEW encryption key:
 - Re-enter NEW encryption key:
4. You will then be prompted for passwords for all of the database locations for all of the environments associated with the repository. If a location cannot be reached, you are prompted to skip the location, re-enter the password, or quit the process.

All encrypted passwords in the `connection-sources.xml` file are re-encrypted using the new key. The encryption key values are not displayed on the window.

Change the repository password

If you need to change the password for your repository, use the password manager utility to update the password in the `connection_sources.xml` file. HPE recommends using the Web Console to modify location passwords for all other databases. For more information, see [Change locations and update passwords, on page 85](#).

NOTE: You can use the password manager utility when required for automation purposes.

To change the password for the repository owner on the database

1. Ensure the encryption key has not been changed.

The master encryption key is required for verification purposes. You can change the encryption key after completing the steps in this section.
2. Open a command window.
3. Run the password manager utility using the following syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <install_directory>/obt/bin/runPwManager.sh -p |
| DOS | <install_directory>/obt/bin/runPwManager.bat -p |

where <install_directory> is the location where you installed the software.

4. Enter values for the prompts. Enter a question mark (?) to see a list of available values.
 - Enter environment to modify:
 - Enter master encryption key:
 - Re-enter master encryption key:

- Enter location name to modify:

TIP: You should enter OBTREP_DB for location.

- Enter NEW password:
- Re-enter NEW password:

The password is changed, and the values saved to `connection-sources.xml` in `config` under the home directory. See [Manage the home directory, on page 39](#).

Update database location passwords

When you create a standard environment, HPE Structured Data Manager creates the necessary schema or database depending upon your underlying database (Oracle, SQL Server, Sybase, or DB2). At the same time, it assigns passwords to those newly created objects. These objects and their passwords are required for HPE Structured Data Manager to perform its various operations on your source and target databases.

NOTE: For non-intrusive environments, HPE Structured Data Manager creates no schema or databases, and therefore requires no passwords for them.

When you create a new, standard environment in Oracle, SQL Server, or Sybase, HPE Structured Data Manager creates the necessary schema or database, and assigns the passwords specified by you. When you create a new environment in DB2, you must provide the necessary operating system user password required to perform operations against the database.

HPE Structured Data Manager stores the current passwords for the schema or databases at the time of standard environment creation because it needs them to perform its operations.

Other than the initial setting of the schema or database passwords for Oracle, SQL Server, and Sybase, database location passwords are managed and changed through the facilities of the underlying database, not HPE Structured Data Manager. Hence, when a schema or database password is changed on the database side, you must update the password stored by HPE Structured Data Manager. For example, suppose that your company requires all schema passwords be changed every other month. When your OBT_IF schema password changes in your Oracle database, you must update the OBT_IF password that HPE Structured Data Manager stores. Otherwise, it will not be able to perform operations against that location.

You can update database location passwords in one of two ways:

- From the Web Console. For more information, see [Change locations and update passwords, on page 85](#).
- From the command line using the password manager utility:

```
runPwManager.bat|sh -p <newpasswd> -l <location> -e <env> -k <encryption>
```

where:

<location> is one of the following:

- OLTP_DB is the source database.
- OBTINTF_DB is the interface schema (OBT_IF).
- ARCHIVE_ACCESS_DB is the archive access schema (OBT_BUILD_AA).

- HISTORY_DB is the target database.
- HISTRELOC_DB is the relocation schema (OBT_RELOC).

<env> is the environment for which you are changing passwords.

<encryption> is the encryption key.

Appendix D: Scripted product setup and business flow deployment

If you plan to set up and deploy the HPE Structured Data Manager multiple times, you can use scripted deployment.

- [Understanding properties files, below](#)
- [Script the repository setup, on the next page](#)
- [Script the environment creation, on page 193](#)
- [Script business flow deployment, on page 200](#)

See the *HPE Structured Data Manager Installation Guide* for information about scripted installation of the product.

Understanding properties files

Scripted deployment requires that you create the following properties files to define the installation.

| Property file | Description |
|---------------------------------|---|
| repository.install.properties | Contains properties to create your repository. |
| product.install.properties | Contains properties to create environments with database to file and optionally database to database. |
| businessflow.install.properties | Contains properties to deploy your business flows. |

You can create these files in one of two ways:

- Use the Web Console to generate the properties files.
- Use a text editor to manually create the files.

The instructions in this chapter assume that you have used the Web Console at least once to create the properties files in the install directory.

NOTE: Properties files should contain property values with ISO 8859-1 character encoding only. For more information on using characters that cannot be expressed using ISO 8859-1 encoding, see [Use special characters in properties files, below](#).

Use special characters in properties files

Properties files should contain property values with ISO 8859-1 character encoding only. Unicode escapes can be used for characters that cannot be directly represented in this encoding. One u character can be used per escape sequence.

To convert properties files that contain non-ISO 8859-1 character encoding

1. Run the Native-to-ASCII Converter, `native2ascii`, with the following syntax:

```
native2ascii -encoding <encoding_name> <input_file> <output_file>
```

| Parameter | Description |
|-----------------|--|
| <encoding_name> | A valid Java encoding. For more information, see the Sun Java website at http://java.sun.com . |
| <input_file> | The original properties file with non-ISO 8859-1 character encoding. |
| <output_file> | The new properties file with only ISO 8859-1 character encoding. |

2. Use the new properties file for your scripted installation.

Script the repository setup

To script your repository installation, you need to create a file with the properties you want to use, and then run the script. The following sections contain the scripting procedure, the properties in the properties file, and examples of the properties file.

NOTE: If you are archiving data from Sybase, you must install your repository on either an Oracle, SQL Server, or PostgreSQL database.

Scripted repository installation procedure

If you choose to include passwords or encryption.key properties in the properties file, they are saved in the text file and can be accessed. To ensure that the passwords are not saved, omit the entire password property line from the properties file. You will then be prompted for the passwords.

1. Make sure that the HPE Structured Data Manager has been installed. See the *HPE Structured Data Manager Installation Guide*.
2. Navigate to the directory that contains the installation scripts:

Example

```
<install_directory>/obt/install
```

where <install_directory> is the location you installed the software.

3. Open the `repository.install.properties` file in a text editor, or create the file if it does not exist.
4. Edit the properties file as necessary for your installation. See [Scripted environment creation properties, on page 193](#).
5. Open a command window.
6. Install the repository using the appropriate syntax:

| Platform | Syntax |
|----------|--|
| UNIX | <code>./deploy_repository.sh <properties_file_name></code> |
| DOS | <code>deploy_repository.bat <properties_file_name></code> |

Where `properties_file_name` is the name of the properties file you are using for the repository.

Scripted repository installation properties

The following properties are used for installing the repository:

| Property name | Expected value |
|--|--|
| <code>encryption.key</code> | The encryption key to be created. If you omit the line, you will be prompted for the encryption key. |
| <code>repository.conn.rdbmstype</code> | sqlserver or oracle or postgresql. |
| <code>repository.conn.dbserver</code> | Database SID or server name for the repository. |
| <code>repository.conn.host</code> | Host name for the repository. |
| <code>repository.conn.port</code> | Port number for the repository. |
| <code>repository.dbadmin.username</code> | Repository database or database server administrator username. |
| <code>repository.dbadmin.password</code> | Repository database or database server administrator password. If you omit the line, you will be prompted for the password. |
| <code>repository.owner.name</code> | Repository user to be created by the installer. |
| <code>repository.owner.password</code> | Repository user password to be created. If you omit the line, you will be prompted for the password. |
| <code>repository.default.storagelocation</code> | Default tablespace to use when creating repository users. (Oracle only) |
| <code>repository.temp.storagelocation</code> | Temporary tablespace to use when creating repository users. (Oracle only) |
| <code>repository.dbname</code> | Repository database name. (SQL Server or PostgreSQL) |
| <code>repository.default.storagelocation.size</code> | Data device size to be allowed to the repository database in MB. (SQL Server or PostgreSQL) |
| <code>repository.log.storagelocation.size</code> | Log device size to be allocated to the repository database in MB. (SQLServer or PostgreSQL) |
| <code>console.admin.username</code> | The username for the Web Console administrator. |

| Property name | Expected value |
|----------------------------|---|
| console.admin.passwd | The password for the Web Console administrator. |
| console.admin.email | Optional. The email address of the Web Console administrator. |
| console.admin.userRealName | Optional. The real name of the Web Console administrator. |
| console.admin.phone | Optional. The phone number of the Web Console administrator. |
| console.admin.description | Optional. The description of the Web Console administrator. |

Scripted repository installation properties file example

The following examples demonstrate the use of the properties file.

Oracle

```
encryption.key=myEncryptionKey
repository.conn.rdbmstype=oracle
repository.conn.dbserver=orcl1
repository.conn.host=myHost
repository.conn.port=1521
repository.dbadmin.username=system
repository.dbadmin.password=myPassword
repository.owner.name=obt_rep
repository.owner.password=myPassword
repository.default.storagelocation=USERS
repository.temp.storagelocation=TEMP
console.admin.username=admin
console.admin.passwd=admin
```

SQL Server

```
encryption.key=myEncryptionKey
repository.conn.rdbmstype=sqlserver
repository.conn.dbserver=MSOLTP
repository.conn.host=myHost
repository.conn.port=1433
repository.dbadmin.username=sa
repository.dbadmin.password=myPassword
repository.owner.name=obt_rep
repository.owner.password=myPassword
repository.dbname=obt_rep
repository.default.storagelocation.size=50
repository.log.storagelocation.size=50
console.admin.username=admin
console.admin.passwd=admin
```


Script the environment creation

To script your environment creation, you need to create a file with the properties you want to use, then run the script. The following sections contain the scripting procedure, the properties in the properties file, and examples of the properties file.

Each environment requires a separate `product.install.properties` file.

Scripted environment creation procedure

The sample `product.install.properties` file is installed with the HPE Structured Data Manager.

If you choose to include passwords or encryption.key properties in the properties file, they are saved in the text file and can be accessed. To ensure that the passwords are not saved, omit the entire password property line from the properties file. You will then be prompted for the passwords.

NOTE: Database to file archiving is installed with all environments. Database to database archiving is optional.

1. Make sure that the HPE Structured Data Manager has been installed. See the *HPE Structured Data Manager Installation Guide*
2. For non-intrusive environments, place the generic JDBC driver in the directory `<install_directory>/obt/lib/jdbc`, where `<install_directory>` is the location where you installed the software.
3. Navigate to the directory that contains the installation scripts:

Example

`<install_directory>/obt/install`

where `<install_directory>` is the location you installed the software.

4. Open the `product.configure.properties` file in a text editor.
5. Edit the properties file as necessary for your installation. See [Scripted repository installation properties file example, on the previous page](#)
6. Open a command window.
7. Create the environment and deploy the product using the appropriate syntax:

| Platform | Syntax |
|----------|---|
| UNIX | <code>./deploy_product.sh <properties_file_name></code> |
| DOS | <code>deploy_product.bat <properties_file_name></code> |

Where `properties_file_name` is the name of the properties file you are using.

Scripted environment creation properties

The following properties are used for creating the environment and deploying database to file:

| Property name | Expected value |
|---|--|
| encryption.key | Encryption key. If you omit the line, you will be prompted for the encryption key. |
| environment.id | The name you want to use for the environment. |
| environment.description | The optional description of the environment. |
| ea.product.configuration | Specifies whether the environment is standard or non-intrusive. For a non-intrusive environment, set the value to <code>lite</code> . Otherwise, set it to default. If you set this property to <code>lite</code> , the following properties are ignored because they are not applicable to a non-intrusive environment: <ul style="list-style-type: none"> • <code>user.index.location</code> • <code>source.dbadmin.username</code> • <code>source.dbadmin.password</code> • <code>source.dbadmin.osauth</code> • <code>source.interface.default.storagelocation</code> • <code>source.interface.temp.storagelocation</code> • <code>source.interface.default.storagelocation.size</code> • <code>source.interface.log.storagelocation.size</code> |
| source.rdbms.name | Acceptable values are <code>oracle</code> , <code>sqlserver</code> , or <code>sybase</code> . |
| source.conn.dbserver | The name of the Oracle database service, SQL Server or Sybase server name, or DB2 data source for the source database. |
| source.conn.host | The host for the source database. |
| source.conn.port | The port number for the source database. |
| source.dbadmin.username | The source database or database server administrator username. |
| source.dbadmin.password | The source database or database server administrator password. If you omit the line, you will be prompted for the password. |
| source.interface.dbname | The interface database name. (SQL Server, Sybase only) |
| source.interface.default.storagelocation.size | The data device size to be allocated to the interface database in MB. (SQL Server only) |
| source.interface.log.storagelocation.size | The log device size to be allocated to the interface database in MB. (SQL Server only) |

| Property name | Expected value |
|--|---|
| source.interface.default.storagelocation | The location for the data device. (Sybase only) |
| source.interface.log.storagelocation | The location for the log device. (Sybase only) |
| source.interface.owner.name | Interface user to be created by the installer. |
| source.interface.owner.password | Interface user password. If you omit the line, you will be prompted for the password. |
| source.interface.default.storagelocation | The default tablespace to use when creating the interface user for the source database. (Oracle only) |
| source.interface.temp.storagelocation | The temporary tablespace to use when creating the interface user for the source database. (Oracle only) |
| user.index.location | The location for the user index. Possible values are source or repository. |

The following optional properties are used for deploying database to database:

| Archive option | Property name | Expected value |
|----------------|---------------------------------------|--|
| | product.configuration | SI or DI |
| | archiveaccess.owner.name | The archive access user to be created. |
| | archiveaccess.owner.password | The password for the archive access user. If you omit the line, you will be prompted for the password. |
| | archiveaccess.location | The location of the archive access user. Acceptable values are OLTP_DB or HISTORY_DB. |
| | archiveaccess.default.storagelocation | The default tablespace to use when creating the archive access user. (Oracle only) |
| | archiveaccess.temp.storagelocation | The temporary tablespace to use when creating the archive access user. (Oracle only) |
| DI | hist.conn.dbserver | The target database server name or SID. |
| DI | hist.conn.host | The host for the target database. |
| DI | hist.conn.port | The port number for the target database. |
| DI | hist.dbadmin.username | The target database server administrator username. |
| DI | hist.dbadmin.password | The target database server administrator password. If you omit the line, you will be prompted for the password. |

| Archive option | Property name | Expected value |
|----------------|---|--|
| DI | hist.reloc.owner.name | Target relocation user (Oracle) |
| DI | hist.reloc.owner.password | Target relocation password (Oracle). If you omit the line, you will be prompted for the password. |
| DI | hist.reloc.default.storagelocation | The default tablespace for creating the target relocation user. (Oracle only) |
| DI | hist.reloc.temp.storagelocation | The temporary tablespace for creating the target relocation user. (Oracle) |
| DI | hist.reloc.dbname | The target database name. (SQL Server only) |
| DI | hist.reloc.default.storagelocation.size | The data device size to be allocated to the target relocation user in MB. (SQL Server only) |
| DI | hist.reloc.log.storagelocation.size | Log device size to be allocated to the target relocation user in MB. (SQL Server only) |
| DI | histreloc.sourceinterface.dblink.name | The database link name to be created from the target relocation schema to the interface schema. (Oracle only) If GLOBAL_NAMES is set to true, then the database link name must be the same as the source database. |
| DI | sourceinterface.histreloc.dblink.name | The database link name to be created from the interface schema to the target relocation schema. (Oracle only) If GLOBAL_NAMES is set to true, then the database link name must be the same as the target database name. |
| DI | source.tnsname | The source TNS name used to create the database link. (Oracle) |
| DI | hist.tnsname | The target TNS name used to create the database link. (Oracle) |

Scripted environment creation properties file example

The following examples demonstrate the use of the properties file.

SQL Server example

SQL Server with database to file only:

```
#Product Information
encryption.key=myEncryptionKey
environment.id=myEnvironment
environment.description=My Default Environment
```

```
#Source Database Information
source.rdbms.name=sqlserver
source.conn.dbserver=MSOLTP
source.conn.host=myHost
source.conn.port=5001
source.dbadmin.username=sa
source.dbadmin.password=myPassword

#Interface User Information
source.interface.dbname=obt_if
source.interface.owner.name=obt_if
source.interface.owner.password=myPassword
source.interface.default.storagelocation.size=50
source.interface.log.storagelocation.size=25
```

Oracle example

Oracle with database to database enabled for distributed archive:

```
#Product Information
encryption.key=myEncryptionKey
environment.id=myOracleEnv
environment.description=Oracle Environment

#Source Database Information
source.rdbms.name=oraOLTP
source.conn.dbserver=myServer
source.conn.host=myHost
source.conn.port=1521
source.dbadmin.username=system
source.dbadmin.password=myPassword

#Interface User Information
source.interface.owner.name=OBT_IF
source.interface.owner.password=myPassword
source.interface.default.storagelocation=ADM_DATA
source.interface.temp.storagelocation=TEMP

product.configuration=DI

#Archive Access Information
archiveaccess.owner.name=OBT_BUILD_AA
archiveaccess.owner.password=myPassword
archiveaccess.default.storagelocation=ADM_DATA
archiveaccess.temp.storagelocation=TEMP
archiveaccess.location=OLTP_DB

#Relocation User Information
hist.reloc.owner.name=OBT_RELOC
hist.reloc.owner.password=myPassword
```

```
hist.reloc.default.storagelocation=ADM_DATA
hist.reloc.temp.storagelocation=TEMP

#Target Database Information
hist.conn.dbserver=oraHIST
hist.conn.host=myHost
hist.conn.port=1521
hist.dbadmin.username=system
hist.dbadmin.password=myPassword

histreloc.sourceinterface.dblink.name=OBT_HIST_TO_OLTP
sourceinterface.histreloc.dblink.name=OBT_OLTP_TO_HIST

#TNS Names
source.tnsname=oraOLTP
hist.tnsname=oraHIST
```

Sybase example

Sybase with database to file only:

```
#Product Information
encryption.key=myEncryptionKey
environment.id=MySybaseEnv
environment.description=Sybase environment

#Repository Information
rdbms.name=postgresql
repository.dbname=obt_rep
repository.conn.rdbmstype=postgresql
repository.conn.dbserver=xxx
repository.conn.host=localhost
repository.conn.port=1527
repository.dbadmin.username=obt_rep
repository.owner.name=obt_rep

#Source Database Information
source.rdbms.name=sybase
source.conn.dbserver=SYBASE_OLTP125
source.conn.host=myHost
source.conn.port=5000
source.dbadmin.username=sa
source.interface.owner.name=obt_if
source.interface.dbname=obt_if
source.interface.default.storagelocation=my_device
source.interface.log.storagelocation=my_log_device
```

DB2 examples

DB2 with database to database:

```
#Product Information
encryption.key=all4one
environment.id=DefEnvD2DDB2
environment.description=DefEnvDB2
product.configuration=DI
#user.index.location=Source

#Source Database Information
source.rdbms.name=db2
source.conn.dbserver=DB2OLTP
source.conn.host=oltp1
source.conn.port=50000
source.dbadmin.username=DB2ADMIN
source.dbadmin.password=

#Interface User Information
source.interface.owner.name=OBTINTF
source.interface.owner.password=
source.interface.default.schema=OBT_IF
source.interface.default.storagelocation=INTF_DATA
#source.interface.temp.storagelocation=TEMP

#Build Archive Access User Information
archiveaccess.owner.name=OBTAA
archiveaccess.owner.password=
archiveaccess.default.schema=OBT_AA
archiveaccess.default.storagelocation=INTF_DATA
#archiveaccess.temp.storagelocation=TEMP
archiveaccess.location=OLTP_DB
aa.to.remotedb.link=HIST_DB

#Target Database Information
hist.conn.dbserver=DB2HIST
hist.conn.host=oltp1
hist.conn.port=50000
hist.dbadmin.osauth=false
hist.dbadmin.username=DB2ADMIN
hist.dbadmin.password=

#Relocation User Information
hist.reloc.owner.name=OBTRELOC
hist.reloc.owner.password=
hist.reloc.default.schema=OBT_RL
hist.reloc.default.storagelocation=RELOC_DATA
histreloc.sourceinterface.dblink.name=OLTP_DB
sourceinterface.histreloc.dblink.name=HIST_DB

ea.product.configuration=default
```

DB2 with database to file only:

```
#Product Information
encryption.key=all4one
environment.id=DefEnvDB2D2F
environment.description=DefEnvDB2D2F
ea.product.configuration=default
user.index.location=Source

#Source Database Information
source.rdbms.name=db2
source.conn.dbserver=DB2OLTP
source.conn.host=oltp1
source.conn.port=50000
source.dbadmin.username=DB2ADMIN
source.dbadmin.password=

#Interface User Information
source.interface.owner.name=OBTINTF
source.interface.owner.password=
source.interface.default.schema=OBT_IF_D2F
source.interface.default.storageLocation=INTF_DATA
```

Script business flow deployment

To script your business flow deployment, you need to create a file with the properties you want to use, and then run the script. The following sections contain the scripting procedure, the properties that can be used, and an example of the properties file.

Scripted business flow deployment procedure

1. Make sure the repository and environment has been created.

NOTE: If you plan to use eligibility analytics, you must enable eligibility analytics in the model for your business flow before deploying, and the model must contain a pause after the selection step. For more information see the *HPE Database Archiving software Developer's Guide*.

2. Navigate to the directory that contains the generated business flows and the businessflow properties files.

Business flows located in the home directory under the main businessflow directory can be accessed by all environments.

Example

/obt/businessflow

See [Manage the home directory, on page 39](#).

Business flows located in an environment directory can only be accessed by that environment.

Example

/obt/businessflow/<environment_name>

3. Make sure that the business flow properties file has the same name as the generated business flow that you want to deploy.

Example

If the full name of the .businessflow file in the business flow directory is ORDERS_D2F_BF.1_0_0_0.busflow, the properties file should be called ORDERS_D2F_BF.1_0_0_0.properties.

4. Edit the properties file with the appropriate property values from [Scripted business flow deployment properties, below](#).
5. Save the file.
6. Open a command window.
7. Navigate to the following directory:

<install_directory>/obt/install

where <install_directory> is the directory where you installed the software.

8. Deploy the business flow using the appropriate command:

| Platform | Syntax |
|----------|--|
| UNIX | ./deploy_businessflow.sh <environment_name> <businessflow_full_name> |
| DOS | deploy_businessflow.bat <environment_name> <businessflow_full_name> |

where <businessflow_full_name> is the full name of the business flow you want to deploy, and <environment_name> is the name of the environment.

Example

```
deploy_businessflow.bat MyEnvironment ORDERS_D2F_BF.1_0_0_0
```

TIP: If you want to remove the job run history of any existing business flows, add the following command: -Dforce.job.history.drop=true

Scripted business flow deployment properties

| Name | Values |
|------------------------------------|--|
| environment.id | Environment ID |
| encryption.key | Encryption key. If you omit the line, you will be prompted for the encryption key. |
| businessflow.install.configuration | Possible values are: <ul style="list-style-type: none">• PrepareOnly• DeployOnly• PrepareAndDeploy |

| Name | Values |
|---|---|
| rdbms.name | The rdbms.name of the source database. Acceptable values are sqlserver, oracle, sybase, or DB2. |
| source.dbadmin.username | The source database administrative username. |
| source.dbadmin.password | The password for the source database administrative user. |
| <cartridge_name>.hist.default.storagelocation | The default target tablespace to use for the specified cartridge. (Oracle only) |
| <cartridge_name>.hist.temp.storagelocation | Temporary target tablespace to use for the specified cartridge. (Oracle only) |
| <cartridge_name>.archiveaccess.default.storagelocation | Default archive access tablespace to use for the specified cartridge. (Oracle only) |
| <cartridge_name>.archiveaccess.temp.storagelocation | Default archive access tablespace to use for the specified cartridge. (Oracle only) |
| <cartridge_name>.hist.default.storagelocation.size | Data device size to be allowed for the specified cartridge on the target database. (SQL Server, Sybase) |
| <cartridge_name>.hist.log.storagelocation.size | Log device size to be allowed for the specified cartridge on the target database. (SQL Server, Sybase) |
| <cartridge_name>.archiveaccess.default.storagelocation.size | Data device size to be allowed for the specified cartridge on the archive access database. (SQL Server, Sybase) |
| <cartridge_name>.archiveaccess.log.storagelocation.size | Log device size to be allowed for the specified cartridge on the archive access database. (SQL Server, Sybase) |
| archiveaccess.owner.username | Required if the business flow contains a database to database cartridge run on Oracle database. (Oracle only) |
| archiveaccess.owner.password | Required if the business flow contains a database to database cartridge run on Oracle database. (Oracle only) |
| hist.dbadmin.username | Required if the business flow contains a distributed archive database to database cartridge. |
| hist.dbadmin.password | Required if the business flow contains a distributed archive database to database cartridge. |

| Name | Values |
|---------------------------------------|---|
| histreloc.sourceinterface.dblink.name | Name of the history to source interface dblink. (Oracle only) |
| sourceinterface.histreloc.dblink.name | Name of the source interface to history dblink (Oracle only) |
| gen.override | Set to true for 3-tier configuration, or false for 2-tier configuration. If omitted or left blank, defaults to false. Required for business flows containing database to file cartridges. |
| source.location.name | The location of the source database. Required for business flows containing database to file cartridges. |
| oa.deploy.config | The database to file deployment configuration. Accepted values are 2tier or 3tier. If omitted or left blank, defaults to 2tier. Required for business flows containing database to file cartridges. |

Scripted business flow deployment properties file example

The following examples demonstrate the use of the properties file.

Oracle example

Oracle database to database business flow

```
#Business Flow Install Properties
businessflow.install.configuration=PrepareAndDeploy
environment.id=MyEnv
rdbms.name=oracle
source.dbadmin.username=system
product.configuration=SI
ORD.hist.default.storagelocation=USERS
ORD.hist.temp.storagelocation=TEMP
ORD.archiveaccess.default.storagelocation=USERS
ORD.archiveaccess.temp.storagelocation=TEMP
```

Oracle example

Oracle database to file business flow

```
#Business Flow Install Properties
businessflow.install.configuration=PrepareAndDeploy
environment.id=MyEnv
rdbms.name=oracle
```

```
source.dbadmin.username=system
source.location.name=OBTINTF_DB
oa.deploy.config=2tier
gen.override=false
```

Glossary

active database

The database from which you plan to move or copy data. Typically, this database is your online transaction processing (OLTP) or production database. In a two- or three-tiered configuration, the active database resides on tier one and is the source for data movement operations.

active environment

The Web Console views and acts upon only one environment at a time, the active environment. To switch the active environment, you use the Change Active option in the Web Console.

activity

In Designer, a component of a business flow, which is added by using the toolbar. For example, you can add archive and reload activities to your business flow. Note, activities in a business flow are different from what you see at runtime and therefore do not necessarily map directly to what you see in the Web Console.

advanced selection

A method of data selection that discovers all of the interrelated rows from multiple tables and conceptually places them in the same application partition for archiving.

annotation

In Designer, a comment associated with the project, or one of its objects or components. These comments are collected and published in a PDF file when you right click a project or business flow and select Generate Documentation.

application partitioning

The concept of partitioning related rows together during data selection, regardless of whether they are in one or more tables. Application partitioning is unique to HPE Structured Data Manager and contrasts with the more common table partitioning offered by the database management

software, which only groups related rows from one table.

archive data store

The location where the data is to be archived. Can be a separate database, separate space on the same database, or an XML file. In a two-tiered configuration, the archive data store resides on tier two and can be a database or XML. In a three-tiered configuration, the archive data store is a database on tier two and XML on tier three, and is the target for data movement operations.

archive query server

The component that provides SQL access to XML database archives.

Autonomy Consolidated Archive (ACA)

A managed, scalable repository that consolidates electronic communications, attachments, and files, and provides complete control over corporate information assets, facilitating compliance with internal corporate governance policies and procedures as well as externally mandated laws and regulations.

business flow

A series of activities, such as archive operations and scripts, that run in sequence. You build business flows in Designer.

business flow status

The Web Console shows the last run of each business flow. The states are Complete/Error/Running.

cartridge

An instance of model- or schema-based eligibility criteria used to move or copy data from one location to another. Cartridges capture the application and business rules to ensure referential integrity of the data. For any one model in your project, you may have many cartridges that use it.

chaining table

The lower level table in a many-to-one or a many-to-many relationship between higher level and lower level tables in the model hierarchy.

classification

The HPE Content Manager (formerly HP Records Manager) classification to be applied to the data moved by HPE Structured Data Manager. This classification specifies where to place the data when it is ingested by HPE Content Manager. For more information, see the HPE Content Manager documentation.

collection

The configuration of a directory location and file pattern to match a set of archived XML files, thus allowing SQL access to the archived data.

comma separated values (CSV)

A database-to-file output format that stores the data as values separated by commas and a metadata file. Each line in the CSV file corresponds to a row in a table. Within a line, fields are separated by commas, each field belonging to one table column. CSV files provide a simple format that many applications can import.

command

Command files or JavaScript files launched by the Web Console on your behalf with status displays.

condition

In Designer, the way you branch your business flow to run or skip an activity based on some criteria.

configuration parameter

A type of parameter that has its values set by an administrator (someone who has repository privileges from the Web Console) through the administrator interface. Typically, this type of parameter represents values that should be changed very infrequently, perhaps only at deployment time.

console user

The Web Console identifies individual users, who are distinct from database users. The properties for a user are User Name, Full Name, Password, Enabled, Description, Email, Phone, and Privileges.

console user name

The login name associated with a Web Console user.

constraint

A column or a list of columns that enables you to identify rows in the database and relate them to one another.

custom properties

User-created name/value pairs in cartridges and business flows. These values are exposed at runtime as parameters.

customization

A change that an administrator or DBA makes to a project provided by a third party, typically for a packaged application like Oracle PeopleSoft or Oracle E-Business Suite. As long as the customization is allowable by the project, the user can merge the customization into newer revisions of the third party project.

customization mode

A Designer mode that provides visual cues to indicate customizations in the model. In a project with locked files, customization mode is on by default, but you can toggle it on and off from the toolbar in the model editor.

data access cartridge

A cartridge that provides lightweight query access to retired or archived data. Data access cartridges are designed by the archive developer but can be run by business users with no technical expertise.

data masking

The process of replacing private or confidential data during movement with a specified mask. You can choose from pre-defined masks that are part of HPE or create your own mask. A mask may or may not be reversible upon reload from the archive data store.

data movement

The method used by HPE Structured Data Manager to actually move data (transactional, bulk or partitioned for database to database, and copy or archive for database to file).

data transparency

The ability to access archived data through your standard application interfaces for data access. Data transparency enables users to access archived data as though it were still in the active database.

database constraint

A constraint that exists in the database and can be discovered and referenced from Designer.

database to database

A movement in which data goes from an active database to an archive database, or separate tablespaces inside the active database. Typically, the archive database is located on cheaper storage devices.

database to file

A movement in which data goes from an active database to a file (XML, JSON or CSV format), which is offline but still accessible through SQL using the archive query server and a client tool of your choice.

deployment assistant

The user interface component within Designer used to deploy or generate business flows.

description

A technical description created by the developer for her own reference. These descriptions do not appear in the generated PDF file for the cartridge or business flow.

Designer

The user interface component used to develop, test, and deploy your archiving solution. Designer is a powerful graphical development environment for archive solutions.

distributed instance

A configuration option for database-to-database archiving where the data you archive is stored on a separate database from the source or active database.

DRE

See [Dynamic Reasoning Engine \(DRE\)](#).

driving table

A driving object is a root of a model hierarchy. Its relationship to the child tables drives the

selection of transactions.

dynamic list of values

A list of values for a parameter that obtains its members from a SELECT statement that returns identifiers and labels.

dynamic parameter

A type of parameter that has its value set by a Groovy script that runs at deployment time to obtain a value. For example, this type of parameter can supply the type or version of a database or application, which can be obtained programmatically at deployment time.

Dynamic Reasoning Engine (DRE)

A platform technology that uses high performance pattern-matching algorithms to search for content stored in Autonomy repositories. Performs core information operations for contextual analysis and concept extraction, enabling solutions for the categorization, summarization, personalization, hyperlinking, and retrieval of all forms of information.

environment

The source and (optional) target credentials against which you plan to run commands. You can define multiple environments within your installation to identify various source and target databases.

error

One of the ways in which you can interrupt a business flow. Error indicates that the business flow failed for some reason.

exclusive rules

One of the ways in which HPE Structured Data Manager determines whether to include or exclude rows from the archive operation. Exclusive rules require all rows in the constraint table to match for inclusion. Exclusive rules exclude the instance if the condition on any child is false, like STATUS='CLOSED'.

exit

One of the ways in which you can interrupt a business flow. You can exit successfully or with a warning.

export

The way that you save an HPE Structured Data Manager project to an exchange format (.hdp) from the File menu. See also *import*.

export data

The way that a user can send data to CSV format from Preview using the toolbar item.

generate documentation

The process of collecting and grouping all annotations into a PDF file that also describes the business flow or cartridge structure.

history schema

For database-to-database archiving, the schema in the target database where the archived data is stored.

IDOL

See [Intelligent Data Operating Layer \(IDOL\)](#).

import

The way that you transfer projects from exchange format (.hdp) into the Project Navigator. You can also use import to migrate cartridges created in 5.1 to 6.x. See also *export*.

inclusive rules

One of the ways in which HPE Structured Data Manager determines whether to include or exclude rows from the archive operation. Inclusive rules require only one row in the constraint table to match the rule and be included. Inclusive rules include the instance if the condition on any child is true, like `PRODUCT_RECALLED='Y'`.

indexing cartridge

A cartridge that indexes your data for better searching. For example, you might associate an indexing cartridge with a database-to-file archiving cartridge to improve performance when querying the archive data files.

Intelligent Data Operating Layer (IDOL)

An information processing layer that collects indexed data from connectors and stores it in a structure optimized for fast processing and retrieval, integrating unstructured, semi-structured, and structured information from multiple repositories.

interrupt

The way to stop or pause a business flow (pause, error, exit with warning, exit successfully).

local deployment

The generation and deployment of your cartridge or business flow to an environment on your local, Designer client. Deployment files are generated locally and then deployed to the designated, local environment.

lookup table

A table that contains helpful non-transactional information. For example, non-transactional information could be status definitions, or the name of the sales representative.

managed table

A table in the model that is copied and then purged from the active database by a cartridge. Transactional, chaining, and driving tables in a model are all typically managed tables.

model

A model identifies the tables and table relationships representing a business entity or related business entities. A project can have multiple models. Each model contains a driving table and all of its child and descendent tables.

model compatibility

Each model in your project can have one or more dynamic parameters associated with it to verify the compatibility with the target environment. If the compatibility parameter returns false, then the cartridge referencing the model will not deploy or run and throw an error. For example, the script could return false for Oracle 10.2 and true for Oracle 11.1 to indicate that a cartridge referencing the model can only deploy and run against Oracle 11.1.

model-based cartridge

A cartridge that moves data based upon a defined data model with relationships. This type of cartridge is typically used for ongoing archive operations.

non-intrusive environment

In a non-intrusive environment, data is archived without an interface schema and a generic

JDBC driver is used. A non-intrusive environment enables you to copy or archive data from read-only sources, which is especially helpful in cases where the data is associated with older technologies that might not support basic SQL statements or when the database administrator or company policy prohibits write access to the production environment.

OLTP database

The online transaction processing database that typically is your active or source database.

pause

One of the ways in which you can interrupt a business flow. Pausing suspends the business flow while awaiting operator intervention.

Content Manager

Enterprise document and records management software designed to simplify the capture, management, security, and access to information. Content Manager enables organizations to more easily comply with regulations and corporate policies, and it helps secure information from inappropriate access and misuse.

reload

The act of taking data from an archive data store and loading it into the active database.

remote deployment

The generation and deployment of your cartridge or business flow to an environment on a system that is remote from your Designer client. Deployment files are generated locally and then deployed to the designated, remote environment.

repository

The location that holds business flow metadata, product configuration data, and data collected during runtime. The repository can be located on your active database or another logical database.

rule

Qualifications added to the model in order to include or exclude data based on certain criteria. For example, you might add a rule to exclude from archiving any orders that are not yet closed.

runtime parameter

A type of parameter that has its values set by the operator executing the job in Console or on the command line. Typically, this type of parameter represents operational values that tend to change frequently and therefore need to be set each time the job is run.

schema-based cartridge

A cartridge that moves data based upon the database schema rather than a defined data model with relationships. This type of cartridge is typically used for database retirement or the cleanup of orphan tables.

selection

The form of data selection to use (standard or advanced) for choosing data. When deploying a cartridge or adding it to a business flow, you must specify the selection method.

single instance

A configuration option for database-to-database archiving where the data you archive is stored on the same database (Oracle) or the same server (SQL Server) as the source or active database.

source

The location (database) from which you are copying or moving data.

SQL access server

See *archive query server*.

standard selection

A method of data selection that restricts itself to the rows identified by the model. Unlike advanced selection, it does not attempt to traverse related rows across multiple tables.

structured records management

A type of solution that extracts structured data from a source application and moves it into XML format. The XML is then ingested into the corporate records management system for long term management and eventual disposal according to corporate policy.

table use

A database table, view, or synonym that is referenced in Designer, for example, in the model. The same table can be used multiple times in a model. For example, a table could be

appear as a transactional table and a lookup table in the same model.

target

The location (database or XML) to which you are copying or moving data.

tier

A level in your database archiving configuration. You can have two- or three-tiered configurations. In a two-tiered configuration, tier one contains your active database and tier two your archive data store, which can be a database or XML. In a three-tiered configuration, tier one contains your active database, tier two an archive database, and tier three XML.

transactional data movement

Transactional movement uses set-based data movement and is the default method of movement.

transactional table

A table that contains information about the business transaction. For example, a transactional table might contain detailed tax or payment information related to each business transaction.

unique identifiers (UIDs)

A 16 hexadecimal identifier calculated based on the content of a Designer file. This value is used to determine if the user has customized key pieces of a project.

unmanaged table

A table in a model that is copied but not purged from the active database by a cartridge. Lookup tables in a model are typically unmanaged tables.

Vertica

Column-oriented SQL database management software for storing and analyzing structured data. Used to manage large, fast-growing volumes of data and provide fast query performance for data warehouses and other query-intensive applications.

virtual constraint

A constraint that you define in Designer that only exists within HPE Structured Data Manager.

Web Console

A browser-based interface where you can create and manage your deployment environments, and deploy, run, administer, and monitor your business flows.

Index

.

.out file
 upgrade output file 133

A

accessing
 archive data files 136
 archived data from command line 105
 creating archive access 98
adding
 hints 164
 session variables 164
advanced selection
 eligibility analytics 176
application
 data directory 39
AQS
 overview 136
archive query server
 See AQS 136
archiveaccess.owner.password 202
archived data
 accessing from command line 105
 creating access 98
 file location 39
 relocating 99
 relocating from command line 106

B

before you begin
 related documentation 14
 tutorials 43, 58
business flows
 before you run on command line 101
 cancelling from command line 104
 configuring 83
 deploying 95
 parameters 83
 running from command line 101
 tutorial for continuing 61
 tutorial for running 46, 59
businessflow.install.configuration 201

C

cancelling
 business flows from command line 104
 failed jobs 97
 failed jobs from command line 103
cartridges
 configuring 83
 parameters 83
character encoding 189
commands
 ALTER SESSION 169
 SET 169
compression
 GZIP 139, 141
 XML 139, 141
configuration
 file location 39
configure
 email 33
configuring
 parameters 82
context values
 about 165
copy
 backend to backend 107
 backend to database 107
 backend to file 107
 file to backend 107
 file to database 107
 file to file 107
creating
 indexes 96

D

data
 accessing from command line 105
 archived, creating access 98
 transparency 98
data movement
 configuring a gateway instance 54
 configuring an application instance 55
 configuring two instances for 53
 for database to database 82
 multiple 53
 sending to gateway instance 56
data type mapping file 117
 editing 119
 validating 119
database archiving
 for performance 14
 long term retention 14

- database to database
 - archive access parameters 156
 - core parameters 151
 - distributed instance parameters 159
 - file naming parameters 161
 - method 82
 - PDM parameters 158
 - performance parameters 152
- database to file
 - core parameters 139
 - performance parameters 143
 - validation parameters 146
 - XAM parameters 147
- databases
 - location, destination 108
- delete
 - partition 179
- deleting
 - from backend 107
 - from database 107
 - from file 107
- deploying
 - business flows 95
 - preventing 182
 - properties file for scripted 189
 - scripted 189
 - tutorial 44
- destinations
 - database locations 108

E

- eligibility analytics
 - advanced selection 176
 - tutorial 58
 - tutorial for viewing 60
- email
 - configuration 33
- encryption keys
 - changing 185
 - deployment properties 201
- environment.id 201
- environment_ID 105
- environments
 - creating 78
 - scripted creation 193
 - tutorial 24

F

- file archive
 - settings 84

G

- GZIP compression 139, 141

H

- Hadoop
 - configuration 88
- hist.dbadmin.password 202
- history schema
 - about 122
 - overview 124
 - preparing 124
 - requirements 124
 - script naming conventions 129-130
- home directory
 - about 39
 - changing 39

I

- IDOL
 - settings 34
- indexes
 - creating 96
 - recreating 96
- installing
 - new repository 66
 - repository, scripted 190
 - scripted 189
- ISO 8859-1 189

J

- JDBC
 - URL 80
- jobs
 - failed 97, 103
 - not responding 97, 103
 - recovery script 103
 - suspended 97, 103

L

- LDAP
 - adding users 74
 - implementing 73
 - importing certificate 73
 - updating webCosole.properties 74
- locking
 - lock_repository script 183
 - repositories 182
- log files
 - checking 133
 - location 39

- log4j.properties 72
- logging levels 72
- obt.log 71
- properties 134
- using 71
- viewing 71
- log4j.properties 134

M

- managed tables 122
 - list of 128
- mapping file 117
 - editing 119
 - validating 119
- monitoring
 - log files 71
- multiple nodes
 - Web Console 70

N

- nodes
 - adding Web Console 70

O

- obt.log
 - main log file 71
 - upgrading log 134
- Oracle
 - partition delete 179
 - RAC 81
 - RAC connection through JDBC 80
 - SQL trace 180
- overwrite
 - script option 129

P

- parameters
 - cartridge 83
 - configuration, database to database 150
 - configuration, database to file 139
 - configuring 82
- partition
 - deletion 179
- passwords
 - changing encryption keys 185
 - changing for repositories 185
 - changing repository passwords 185
- patch numbers
 - most recent 125

- ports
 - starting Web Console on 65
- prerequisites
 - product 10
- program data
 - directory 39
- purge
 - partition 179

R

- reclaiming
 - space 177
- recovering
 - failed jobs 97
 - failed jobs from command line 103
 - script for failed jobs 103
- reloading
 - before you begin 96
 - tutorial 47
- relocating
 - archived data 99
 - archived data from command line 106
- repositories
 - changing passwords 185
 - installing 66
 - lock_repository script 183
 - locking 182
 - scripted setup 190
 - tutorial for installing 19
 - unlock_repository script 184
 - unlocking 182, 184
 - upgrading 69
- repository
 - querying 128
- running
 - before you begin on command line 101
 - business flows from command line 101
- runPwManager 185-186

S

- schema mapping
 - simple example 50
- schema mapping file 115
- script
 - deploying 132
 - editing 129
 - generating 128
 - location 130
 - verifying results 133
- seed statements
 - about 165

- files 165
- session variables
 - about 169
- settings
 - file archive 84
- source and target
 - discovering differences 127
- source.dbadmin.password 202
- SQL
 - hints 164
 - trace for Oracle 180
 - tuning 164
 - tuning hints 164
 - tuning properties 164
- sql_tuning.properties
 - default hints 171
- starting
 - Web Console 63
 - Web Console on different port 65
 - Web Console tutorial 17
- stopping
 - Web Console 63, 65

T

- target and source
 - discovering differences 127
- TODO comments
 - upgrade script 130
- transparency
 - archived data 98
- tutorials
 - before you begin 43, 58
 - configuring Web Console 17
 - continuing paused business flow 61
 - deploying business flows 44
 - eligibility analytics 58
 - reloading 47
 - running business flow 46, 59
 - undo business flows 47

U

- undo
 - tutorial for running 47
- unlocking
 - repositories 184
 - unlock_repository script 184
- update
 - IDOL settings 34
- upgrading
 - checkonly mode 127
 - confirming need for 126

- deploying script 132
- editing script 129
- history schema about 122
- history schema overview 124
- history schema requirements 124
- log files 133
- patch numbers 125
- preparing for history schema upgrade 124
- repository 69
- test environment 124
- verifying results 133
- XML script 128
- user
 - ARCHADMIN 181
- users
 - tutorial on creating 40

V

- Vertica
 - accessing 52
 - creating a location 50
 - documentation 52
 - introduction 49
 - schema mapping 50
 - upload considerations 49

W

- Web Console
 - adding nodes 70
 - creating environments 78
 - deploying business flows 95
 - installing repository 66
 - installing repository tutorial 19
 - overview 15
 - starting 17, 63
 - starting on different port 65
 - stopping 63, 65
 - tutorial for configuring 17
 - tutorial for environments 24
 - tutorial on creating users 40
 - upgrading repository 69

X

- XML
 - compression 139, 141
- XML archive script
 - COPY_BE_TO_BE 107
 - COPY_BE_TO_DB 107
 - DELETE_FROM_BE 107
 - DELETE_FROM_DB 107
 - MOVE_BE_TO_BE 107

MOVE_BE_TO_DB 107

Send documentation feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Runtime Guide (HPE Structured Data Manager 7.52)

Add your feedback to the email and click **Send**.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to AutonomyTPFeedback@hpe.com.

We appreciate your feedback!

