



Operations Orchestration

Software Version: 10.80

Windows and Linux

Administer

Document Release Date: September 2017

Software Release Date: September 2017



Hewlett Packard
Enterprise

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise 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 Hewlett Packard Enterprise 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

© September 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

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

This site requires that you register for an HP Passport and to sign in. To register for an HP Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** on the HP Passport login page.

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

Contents

Administrating Operations Orchestration	10
Initial Administration Tasks	10
Configuring and Enabling Authentication	10
Configuring Authentication	10
Enable Authentication	11
Setting the System Locale in central-wrapper.conf	11
Managing OO Licenses	11
Viewing License Information	12
Installing a New License	13
Purchasing a License	13
Change the MS SQL Collation Language or the central-wrapper.conf Language for Content	14
Adjusting the Logging Levels	14
Configuring LWSSO Settings	15
Changing the Service Name and Description	17
Changing the Timeout Limit for the Operations OrchestrationWeb Interface	18
Scheduling Database Purging	18
Purging Flows	18
Purging APIs	19
Specify Communication Protocols	19
Recurring Administration Tasks	20
Encrypting and Obfuscating Passwords	20
Changing the Database Password	22
Changing the Database IP	22
Changing the Hostname/IP of Central or RAS	23
Adjusting the Timing of Quartz Jobs	24
Changing the URL of a Central/Load Balancer on the RAS Side	25
Backing Up Operations Orchestration	25
Backing Up	25
Recovery	26

Setting up Disaster Recovery	26
Planning for Disaster Recovery	26
Performing Disaster Recovery	26
System Tuning	27
Database Maintenance	28
Set Up Database Environment	28
Introduction to Preparing the Database Environment	28
Overview	28
Using Database Clusters	30
Database Security	30
Database Schema	31
OO Database Sizing	31
Hardware Requirements	33
Deploying and Maintaining a Microsoft SQL Server Database	33
Workflow for Microsoft SQL Server Deployment	33
System Requirements for Microsoft SQL Server	34
Hardware Requirements	34
Software Requirements	34
Examples of Tested Deployments	35
Language Support	35
Configuring SQL Server	37
Manually Creating an OO Database on Microsoft SQL Server	38
Manually Creating Database Objects	40
Microsoft SQL Server Database Maintenance	40
Backing up the OO Database	40
Creating a Maintenance Plan	41
Upgrade/Rollback Instructions	44
Database Server Upgrade from 2012 or 2008R2 to Server 2014 or Server 2016	44
Always On Support	45
Deploying and Maintaining an Oracle Database	46
Workflow for Oracle Deployment	46
System Requirements for Oracle	47
Hardware Requirements	47
Software Requirements	47

Oracle Connector	48
Examples of Tested Deployments	48
Language Support	49
Configuring an Oracle Database	49
Manually Creating an OO Database on an Oracle Instance	51
Manually Creating Database Objects	52
Connecting to Oracle Using SID or Service Name	52
Oracle Database Maintenance	53
Backing up the Operations Orchestration Database	53
Creating a Maintenance Plan	53
Upgrade/Rollback Instructions	56
Deploying and Maintaining a MySQL Database	57
Workflow for MySQL Deployment	57
System Requirements for MySQL	57
Hardware Requirements	58
Software Requirements	58
MySQL Connector	59
Examples of Tested Deployments	59
Language Support	59
Configuring MySQL	60
Manually Creating an OO Database on MySQL	62
Manually Creating Database Objects	64
MySQL Database Maintenance	64
Backing up the OO Database	65
Creating a Maintenance Plan	65
Utility for Purging Historical Data	66
Upgrade/Rollback Instructions	66
Deploying and Maintaining a Postgres Database	66
Workflow for Postgres Deployment	67
System Requirements for Postgres	67
Hardware Requirements	67
Software Requirements	67
Examples of Tested Deployments	68
Language Support	69
Configuring Postgres	69

Manually Creating an OO Database on Postgres	70
Manually Creating Database Objects	71
Postgres Database Maintenance	71
Backing up the OO Database	72
Creating a Maintenance Plan	72
Utility for Purging Historical Data	73
Upgrade/Rollback Instructions	74
Database Purging	74
Purge Flows	74
Purge Best Practices	75
Additional Guidelines for Microsoft SQL Server	76
Using Windows Authentication to Access Microsoft SQL Server Databases	76
Configuring OO to Work with Windows Authentication	76
Configuring OO to Work with Always On	77
Additional Guidelines for Oracle	79
Oracle Real Application Cluster (RAC)	79
Single Client Access Name	80
Configuring OO to Work with Oracle RAC	80
Installation Wizard 'Other database' Option	81
Microsoft SQL Server Named Instance Example	82
Microsoft SQL Server Windows Authentication Example	84
Oracle RAC Example	86
Securing and Hardening OO	86
Introduction	87
Security Overview	88
Security Concepts	89
Secure Implementation and Deployment	92
OO Security Hardening	92
Physical Security	92
Secure Installation Guidelines	92
Supported Operating Systems	92
Operating System Hardening Recommendations	92
Tomcat Hardening	93
Installation Permissions	93

Network and Communication Security	93
Communication Channel Security	94
Administration Interface Security	94
Accessing the Administration Interface	94
Securing the Administration Interface - Recommendations	95
User Management and Authentication	95
Authentication Model	95
Types of Users	96
Authentication Administration and Configuration	96
Database Authentication	96
Authorization	96
Authorization Model	96
Authorization Configuration	97
Configure OO Central in IDM Mode	98
Switch to the IDM security profile	98
Update the idm.properties file	98
Configure the return URL of Central in IDM mode	99
Example of the idm.properties file	100
Hidden Tabs in the IDM Mode	100
Roles in IDM	101
Configure LDAP in IDM Mode	102
Add an LDAP configuration for the IDM server	102
Backup	104
Encryption	104
Encryption Model	104
Encryption Administration	105
Digital Certificates	106
Sensitive Information in a Content Pack	108
Hardening OO	108
Security Hardening Recommendations	108
Working with Server and Client Certificates	109
Encrypting the Communication Using a Server Certificate	110
Replacing the Central TLS Server Certificate	110
Importing a CA Root Certificate to the Central TrustStore	112
Importing a CA Root Certificate to a RAS TrustStore	112

Importing a CA Root Certificate to the OOSH TrustStore	113
Importing a CA Root Certificate to the Studio TrustStore	114
Checking the Revoke Status of a Certificate	116
Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password	117
Changing the KeyStore, TrustStore, and Server Certificate Passwords in the Central Configuration	117
Changing the RAS, OOSH, and Studio TrustStore Passwords	119
Encrypting and Obfuscating Passwords	119
Removing Vulnerable Ciphers from the SSL supported Ciphers ..	121
Changing the HTTP/HTTPS Ports or Disabling the HTTP Port	121
Changing Port Values	122
Disabling the HTTP Port	122
Troubleshooting the HTTPS Connector	123
Client Certificate Authentication (Mutual Authentication)	123
Configuring Client Certificate Authentication in Central	124
Updating the Configuration of a Client Certificate in RAS	126
Configuring a Client Certificate in Studio Remote Debugger	127
Configuring a Client Certificate in OOSH	128
Processing Certificate Policies	129
Processing a Certificate Principal	129
Enabling _HPc_Basic_Variables._HP_Product_Acronym to Read from the Subject Alternative Name Field in a Certificate	130
Signing OO Content Packs	130
Verifying the Integrity of the Studio Tools JARs	132
Configuring OO for FIPS 140-2 Level 1 Compliance	134
Prerequisite Steps for Upgraders	136
Configuring Compliance with FIPS 140-2	136
Step 1: Configure the Properties in the Java Security File	137
Step 2: Configure the encryption.properties File and Enable FIPS Mode	138
Step 3: Create FIPS-Compliant Encryption	138
Step 4: Re-encrypt the database password with the new encryption	139
Step 5: Start OO	139

Replacing the FIPS Encryption	139
Changing the FIPS Encryption Key on Central	139
Changing the RAS Encryption Properties	140
Configuring the TLS Protocol	140
Preventing Flows from Accessing the Central/RAS Local File System	141
Adding Java Security Manager	142
Configuring RAS to Add Java Security Manager	142
Configuring Central Embedded Worker to Add Java Security Manager	143

Administering Operations Orchestration

This section documents relevant information regarding the following Operations Orchestration (OO) configuration tasks.

- Completing advanced configuration tasks.
- Configure four different types of database, to be used with OO.
- Deploying and managing OO instances in a secure manner, and configuring security hardening for OO.

Initial Administration Tasks

This section includes administration tasks that you will generally perform once, while setting up Operations Orchestration after installation.

Configuring and Enabling Authentication

After installing Operations Orchestration, you must configure the authentication mechanism and enable authentication.

Configuring Authentication

Select the authentication mechanism to be used in Operations Orchestration, for example, LDAP, SSO, internal users, etc.

For more information on how to configure the authentication mechanism, see "Setting up the System Configuration" in the *OO Central User Guide*.

Enable Authentication

This is done in Central by marking the **Enable Authentication** check box.

Note: The **Enable Authentication** check box is only available if there are existing internal or LDAP users with permission to disable the authentication in the future.

For information about how to enable authentication, see *OO Central User Guide*.

Setting the System Locale in central-wrapper.conf

If your Operations Orchestration (OO) system is localized, you will need to set the properties to reflect the system locale, in the **central-wrapper.conf** file.

1. Open the **central-wrapper.conf** file in a text editor (located under **<installation_folder>/central/conf**).
2. Edit the following properties:

```
set.LANG=  
set.LC_ALL=  
set.LANGUAGE=  
wrapper.java.additional.<x>=-Duser.language=  
wrapper.java.additional.<x>=-Duser.country=
```

For example, for Japanese:

```
set.LANG=ja_JP  
set.LC_ALL=ja_JP
```

3. Save the **central-wrapper.conf** file.


Managing OO Licenses

When Operations Orchestration (OO) is first installed, this installs the Trial version license. You will need to install an Enterprise Edition license within 90 days.

OO offers the following options for licensing:

- **OO Enterprise Edition Trial:** A temporary license that runs for 90 days, and includes all out-of-the-box content packs. After the trial license is expired, you will need to purchase the OO Enterprise Edition license to use the content packs. This is the default license that comes with the OO installer.
- **OO Enterprise Edition:** A full license, unlimited in time, which includes all out-of-the-box content packs. This license unlocks the OO Community Edition and OO Enterprise Edition Trial licenses.
- **OO Community Edition:** A partial license, which includes the Base content pack and the CE content packs, and enables 500 runs per month.
- **OO Enterprise Edition With Limited Functionality:** A partial license that limits all new functionalities that were introduced in OO version 10.80 and later. All features introduced in previous versions of OO are accessible as under OO Enterprise Edition License.

Viewing License Information

You can view the license information by clicking the **Info**  button in the top right corner in Central, to display the **About** popup window. This window displays the current status of the license: if the license is limited in time or number of runs, it shows how much is left.

About HPE Operations Orchestration ×

OPERATIONS ORCHESTRATION


Version:
10.70 - Trial Edition

Build Number:
1368

Licensing:
You have 84 days left in the trial period.

(c) Copyright 2015 Hewlett Packard Enterprise Development Company, L.P.

Open source and third-party software license agreements for this product are available [here](#)

 **Hewlett Packard
Enterprise** Close

Installing a New License

To install a new license, got to the Central **License** tab under **System Configuration > System Settings**.

For more information about setting up the license, see *Central User Guide*.

Purchasing a License

Option 1

Download the license from the OO web site at:

<http://enterpriselicense.hpe.com/redirector/home>

Alternatively, you can go to:

<https://h30580.www3.HPE.COM/poeticWeb/portalintegration/hppWelcome.htm>

After logging in, enter the EON (Entitlement Order Number) in the Welcome page to see your license.

Option 2

Issue a license using the License Management system:

- For a standalone Central installation, issue the license with the IP address of the Central server.
- If you have CSA and OO installed on the same machine, issue one license for that IP address.
- If you have CSA and OO installed on separate machines, issue two licenses for the two IP addresses.
- For a cluster, choose one of the nodes and issue a license for the IP address of this node.

Note: Make sure to open the Central UI of the selected node (and not via the Load Balancer IP) when you install the license.

- If you have CSA and an OO cluster, issue two licenses: one for CSA and one for a selected node.

Note: Make sure to open the Central UI of the selected node (and not via the Load Balancer IP) when you install the license for OO.

For more information about licensing details, contact your Sales team or Account Manager.

Licensing a Cluster

For a cluster, make sure that you issue the license on one of the cluster IPs (and not on the load balancer IP). Apply it directly on the specific node. This sets the license for the entire cluster.

Change the MS SQL Collation Language or the `central-wrapper.conf` Language for Content

When Operations Orchestration (OO) Central was installed, there was an option to select a supported language in addition to English. This language is used as the MS SQL collation language, if relevant, and for content. This language support is required if, for example, you need to ping a server that is configured in Japanese. You can change this language by editing the **central-wrapper.conf** file.

1. Open the **central-wrapper.conf** file in a text editor (located under `<installation_folder>/central/conf`).
2. Edit the language property. For example, change `wrapper.lang=en_US` to `wrapper.lang=ja_JP`.
3. Save the **central-wrapper.conf** file.

For more information about selecting the MS SQL Server database collation language, see "Language Support" in the OO *Database Guide*.

Adjusting the Logging Levels

Because logging activity can slow the performance of Operations Orchestration (OO) and create very large log files, it is important for OO to run with appropriate logging levels. The default logging level is set to INFO. This level provides the necessary information without impacting performance.

In addition, you can adjust the granularity of the information in the log, separately for regular logging, deployment, and execution.

The granularity options are:

- INFO - Default logging information
- DEBUG - More logging information

- WARN - Less logging information
- ERROR - Less logging information

To adjust the granularity in the logging:

1. Open the **log4j.properties** file (under **<installation_folder>/central/conf/log4j.properties**).

Replace INFO with DEBUG, WARN or ERROR in the relevant places in the **log4j.properties** file.

For example:

```
log.level=INFO
execution.log.level=DEBUG
deployment.log.level=DEBUG
```

Note: The Tomcat server keeps access logs under **<installation_folder>/central/var/logs**. To prevent these log files from growing in size, it is recommended that you monitor the size of the files named **localhost_access_log.yyyy-mm-dd.txt**, and delete old files as necessary. Use established best practices for your operating system, whether Windows or Linux.

For example, to delete files older than 10 days:

Unix

```
find <OO_HOME>/central/var/logs/ -name " localhost_access*.txt"
-type f -mtime +10 -exec rm -f {} \;
```

Windows

```
C:\Users\Administrator>forfiles /p "<OO_HOME>\central\var\logs" /s /m
localhost_access*.txt /d -10 /c "cmd /c del @PATH"
```

where **<OO_HOME>** refers to the location of the OO installation folder

You can set this up using a schedule in OO or use Cron or Windows Scheduler to run these commands the background.

Configuring LWSSO Settings

When you install (OO) 10.10 or later, if you choose to upgrade the LWSSO settings from version 9.x, these LWSSO settings will be migrated, but LWSSO will be disabled in OO 10.x, even if it was

previously enabled in version 9.x.

When you enable LWSSO afterward, you may receive warnings under certain scenarios. To clear the warnings from the log, follow the steps below to set the management URL property using the fully qualified domain name.

Central and RAS Installed on the Same Machine

When Central and a RAS are installed on the same machine, and the LWSSO settings are enabled, you must set the management URL property using the fully qualified domain name.

1. Stop the RAS process.
2. In the **ras/conf/ras-wrapper.conf** file, change

```
wrapper.java.additional.<x>=-Dmgmt.url=<protocol>://localhost:<port>/oo  
to
```

```
wrapper.java.additional.<x>=-  
Dmgmt.url=<protocol>://<FullyQualifiedDomainName>:<port>/oo
```

3. Start the RAS process.

RAS Installed on a Different Machine

When the RAS is installed on a different machine from the Central and LWSSO settings are enabled, you must specify the management URL of the Central with the fully qualified domain name during the RAS installation, instead of the IP address.

Connecting Another Application to Central Through LWSSO

When connecting another application to Central through LWSSO, you must specify the management URL of the Central with the fully qualified domain name.

1. Stop the Central process.
2. In the **central/conf/central-wrapper.conf** file, change

```
wrapper.java.additional.<x>=-Dmgmt.url=<protocol>://localhost:<port>/oo  
to
```

```
wrapper.java.additional.<x>=-  
Dmgmt.url=<protocol>://<FullyQualifiedDomainName>:<port>/oo
```

3. Start the Central process.

For information about setting up LWSSO in Central, see *Central User Guide*.

Changing the Service Name and Description

If you want to change the service name and description from the default values, you can edit these in the **central-wrapper.conf** file.

Note: This procedure can be applied to a RAS as well. Just replace the word "Central" with "RAS". For example, use **ras-wrapper.conf**, **ras\conf**, and so on.

1. Stop the Central service.
2. Navigate to the **central\bin** directory. The default is **C:\Program Files\Hewlett-Packard Enterprise\ Operations Orchestration**.
3. Run **uninstall-central-service.bat**.
4. Verify that the service is uninstalled.
5. Navigate to the **central\conf** directory.
6. Make a copy of the **central-wrapper.conf** file.

Note: This step is an optional step used for saving the original configuration, so that you have a backup if a rollback is required or if the modified file becomes unusable in any way.

7. Edit the **central-wrapper.conf** file, by modifying the following lines to suit your needs:

```
wrapper.name  
wrapper.displayname  
wrapper.description
```

For example:

```
# Name of the service  
wrapper.name=00Central  
  
# Display name of the service  
wrapper.displayname= Operations Orchestration Central  
  
# Description of the service  
  
wrapper.description=Operations Orchestration Central is the run time  
environment of Operations Orchestration. It is used for running flows,  
monitoring the various runs, and generating reports.
```

8. Run **install-central-as-service.bat**.

9. Verify that the service is installed with the new names and description.
10. Start the new Central service.

Changing the Timeout Limit for the Operations OrchestrationWeb Interface

You can change the default session timeout (30 minutes) of the Central server by configuring the **tomcat web.xml** file.

1. Open the `<installation_folder>/central/tomcat/conf/web.xml` file with any XML editor.
2. Search for the following:

```
<!-- ===== Default Session Configuration ===== -->  
<!-- You can set the default session timeout (in minutes) for all newly -->  
<!-- created sessions by modifying the value below. -->  
  
<session-config>  
    <session-timeout>30</session-timeout>  
</session-config>
```

3. Change the **session-timeout** value (in minutes).
4. Restart the Central server.

Scheduling Database Purging

In order to keep the Operations Orchestration (OO) database size under control, it is important to schedule database purging operations from day one.

Purging Flows

The purging flows are available in the **HPE Solutions** content pack (available on [ITOM Marketplace](#)). It is recommended to deploy this content pack, configure the flows with your required settings, and schedule them in Central.

The following purging flows are located under **Library > Integrations > Hewlett-Packard Enterprise > Operations Orchestration > 10.x > Database**:

- **Purge Execution Summary** - Purges run data.
It is recommended to schedule this flow as soon as you start to run flows.
- **Purge Debug Events** - Purges Studio Remote Debugger event data.
If you use the Studio Remote Debugger, schedule this flow.
- **Purge Audit Records** - Purges old auditing records if auditing is enabled.
If you have enabled security auditing, schedule this flow.
- **Purge Rerun Info** - Purges rerun data.
If you have defined rerun points in your flows, schedule this flow.

For more information about these purging flows, see the flow descriptions in Central.

Purging APIs

As an alternative to using the purging flows, you can perform purging functions via API.

The following APIs are available:

- **DELETE /audit/records**: Purges old auditing records if auditing is enabled.
- **DELETE /debugger-events**: Purges Studio Remote Debugger event data.
- **DELETE /executions**: Purges run data such as bound inputs, outputs and step log events. This run data needs to be purged regularly, because running large numbers of runs can cause the database to reach the maximum table size.

Note: This only affects the data of completed runs.

- **DELETE /executions/rerun**: Purges the rerun data from the database.
- **DELETE /steps-log**: Purges step data according to time and number of executions to purge.

By using the purge APIs, you can purge the data manually as required, or by scheduling recurrent flows that incorporate these APIs.

Specify Communication Protocols

To specify the communication protocols that must be used by the system notification feature, add the `system.notifications.smtp.enabled.protocols` property in the `central-wrapper.conf` file. To do this:

1. Open the **central-wrapper.conf** file located in the `<installation_directory>\central\conf` directory.
2. Add the `system.notifications.smtp.enabled.protocols` property as follows:

```
wrapper.java.additional.32=-DCentral.login.suspend=300000  
wrapper.java.additional.33=-Doverride.startup.mode=LWSSO  
wrapper.java.additional.34=-Dsystem.notifications.smtp.enabled.protocols=TLSv1.2 TLSv1.1  
  
# Initial Java Heap Size (in MB)  
wrapper.java.initmemory=512
```

The `system.notifications.smtp.enabled.protocols` property is optional and its default value is `TLSv1.2 TLSv1.1 TLSv1 SSLv3`. You can overwrite the default value and specify the protocol that you want as a list of space-separated values.

For example, to specify more than one protocol:

```
wrapper.java.additional.34=-Dsystem.notifications.smtp.enabled.protocols=<protocol 1> <protocol 2> <protocol 3>
```

Recurring Administration Tasks

This section includes administration tasks that you may need to perform periodically or in response to a situation that occurs.

Encrypting and Obfuscating Passwords

You can encrypt or obfuscate a password using the `encrypt-password` script, which is located in `<installation_folder>/central/bin`.

Our recommendation is to use encryption.

Important! After using the `encrypt-password` script, clear the command history.

This is because, on a Linux OS, the password parameter will be stored in cleartext under `/$USER/.bash_history` and accessible by the `history` command.

Encrypting Passwords

1. Locate the `encrypt-password` script, in `<installation_folder>/central/bin`.
2. Run the script with the `-e -p <password>` option, where `password` is the password you want to encrypt.

Note: You can either use `-p` as the flag to encrypt the password or `--password`.

The encrypted password should appear as follows:

```
{ENCRYPTED}<some_chars>.
```

Obfuscating Passwords

1. Locate the `encrypt-password` script, in `<installation_folder>/central/bin`.
2. Run the script with the `-o <password>` option, where `password` is the password that you want to obfuscate.

The obfuscated password should appear as follows:

```
{OBFUSCATED}<some_chars>.
```

Creating a Prompt for the Password

It is recommended to run the `encrypt-password` script without providing the `-p` argument. For example:

```
C:\Program Files\Hewlett-Packard\HP Operations Orchestration\central\bin>encrypt-password  
Password (typing will be hidden):  
Confirm password (typing will be hidden):  
{ENCRYPTED}gAkPCLQsYDhoR1Y2q9BjCQ==  
C:\Program Files\Hewlett-Packard\HP Operations Orchestration\central\bin>
```

This will create a prompt for the hidden password inputs.

Changing the Database Password

1. If Central is up and running, stop the Central Service.
2. Run the `encrypt-password` script with the `--password <password>` option, where `<password>` is the database password.

Note: You can either use `-p` as the flag to encrypt the password or `--password`.

For example:

```
<install-dir>/central/bin/encrypt-password --password <plain-text-pass>
```

3. Copy the result. It should appear as follows:

```
${ENCRYPTED}<some_chars>.
```

4. Go to the folder `<installation_folder>/central/conf`, and open the `database.properties` file.

Change the `db.password` value to the one that you copied. In a cluster, make sure to do this on each node.

The password is stored encrypted, so make sure you include the `{ENCRYPTED}` prefix. The `=` character needs to be escaped. For example, `db.password=GFDGDS12\=\=`.

Changing the Database IP

This section is relevant when you need to configure Operations Orchestration to work with another database instance. All the database parameters, such as database credentials, schema name, tables, and so on, should be identical.

1. Edit the file `<installation_folder>/central/conf/database.properties`.
2. Look for the `jdbc.url` parameter. For example:

```
jdbc.url=jdbc\:jtds\:sqlserver\://16.60.185.109\:1433/schemaName;sendStringParametersAsUnicode\=true
```

3. Change the IP address\FQDN of the database server.
4. Save the file.
5. Restart Central.

Changing the Hostname/IP of Central or RAS

If the Central hostname or IP has changed:

If the hostname/IP has changed in a Central, you will need to manually update this hostname/IP address in a number of places.

1. Edit the file `<installation_folder>/central/conf/central-wrapper.conf`.
2. Look for the `-Dmgmt.url` parameter and update it with the correct hostname or IP address. For example:

```
-Dmgmt.url=http://[machine_name or IP]:8080/oo
```

3. Update the same `-Dmgmt.url` parameter in the file `<installation_folder>/ras/conf/ras-wrapper.conf`.

You need to perform this step for each RAS that communicates with the Central.

4. If you are using a Central certificate, generate a new Central certificate with the new FQDN.

It is not recommended to use the IP.

You will also need to replace the certificate in the Central keystore.

5. If there is a load balancer link configured, update this link in Central.

External URL

URL:

`https://my.server.com:443/oo`

URL of the load balancer, reverse proxy, or DNS load balancer

6. Restart Central and all affected RASes.
7. If you have configured a remote debugging session that uses the Central hostname or IP address, you will need to modify the connection URL accordingly in Studio.

For more information, see the User Guide.

If the RAS IP has changed:

If the RAS IP is changed, this change affects the presentation in the Central URL, under **System Configuration > Topology > Workers**.

- If the RAS is a “reverse RAS”, it will become disabled in the **Topology > Workers** tab. You need to edit the row, and replace the previous IP with the new one.
- If the RAS is a “standard RAS”, the change will not affect functionality, because it is the RAS that connects to the Central.

Adjusting the Timing of Quartz Jobs

In the Operations Orchestration system, quartz jobs run periodically for maintenance of the system.

Each job runs for a set amount of time, and is repeated at set intervals. The following are examples of job triggers:

Trigger Name	Current Repeat Interval	What Happens
onRolling:OO_EXECUTION_STATES_Trigger	4.5 minutes	Rolling the state tables for purging
queueCleanerTrigger	1 minute	Purging the queue tables
queueRecoveryTrigger	2 minutes	Checks if the system needs recovery
recoveryVersionTrigger	0.5 minute	Version counter to be used for the recovery
splitJoinTrigger	1 second	Joins finished splits
onRolling:OO_EXECUTION_EVENTS_Trigger	12 hours	Rolling the events table for purging
Note: This trigger is deprecated.		

If you want to tweak these job timings in order to improve performance, perform the following:

Note: Any change to the timings can have a major effect on the system. Consult with your HPE service representative before making any changes to these triggers.

1. Open the OO tab. Under **MBeans**, there is an operation named **jobTriggersMBean**.
2. Use this operation, and enter the values on the right tab, using the name of the trigger you want to change. Use the exact same name as the table, with a new value for the repeat interval.

This changes the triggering times of the job.

Note: The events persistency mechanism is deprecated (see **onRolling:OO_EXECUTION_EVENTS_Trigger**). You can configure this job if you use the Remote Debugger.

Changing the URL of a Central/Load Balancer on the RAS Side

While the best practice is to configure the URL of a Central/load balancer via the installer, if you need to make changes to the URL after the RAS was already installed, you can do this by editing the **ras-wrapper.conf** file.

For example, this would be needed if you installed a RAS against a Central/load balancer and the Central/load balancer's FQDN changed. You will need to change Central/load balancer's URL stored at the RAS level, so that the RAS is able to communicate again with the Central/load balancer.

1. Stop the RAS.
2. Open the **ras-wrapper.conf** file, located under **<installation folder>/ras/conf**.
3. Edit the URL in the following line:

```
wrapper.java.additional.<x>=-Dmgmt.url=<http or https>://<FQDN>:<port>/oo
```

4. Restart the RAS.

Backing Up Operations Orchestration

Backing Up

Back up the entire content of the **<installation_folder>central\var\security** folder and the **<installation_folder>central\conf\database.properties** file.

Recovery

1. Install a new Central with an existing schema. The installation will fail on the **Start Central** step.
2. Stop the Central service and verify that Central is not running.
3. Open the **central/var/security** folder.
4. Delete the **credentials.store** file, if it exists.
5. Override the **encryption_repository** and **encryption.properties** files with the backed up versions.
6. Override the **central/conf/database.properties** file with the backed up file.
7. Start the Central service.

Setting up Disaster Recovery

Planning for Disaster Recovery

To enable disaster recovery, create a secondary datacenter by replicating the following:

- Database configuration
- Schema
- OO configuration and security files: **<installation_folder>\central\var\security** and **<installation_folder>\central\conf**

The replication can be ongoing, using whatever means you choose, for example, disk replication or database replication.

Performing Disaster Recovery

OO supports cold disaster recovery, which requires a manual process to switch over from the primary datacenter to the secondary datacenter. This may be required in the case of a full or partial failure in the primary datacenter.

1. Restore the OO server(s) by reinstalling all nodes (Central and RAS). Use the database configuration, schema, and OO configuration and security files that you replicated for the secondary datacenter.

Note: You can configure the database (by editing the **database.properties** file). See the *OO Administration Guide* for more information about configuring the database connection details.

2. Delete the old worker nodes in the database, which no longer exist.
3. Turn off OO in the primary datacenter.
4. Switch over to the secondary datacenter.

System Tuning

You can configure one or more of the following parameters for each Central node. The JVM and worker threads are also applicable to each RAS node.

Note: Changing any of the parameters below requires a restart of the configured node.

- **JVM heap size** - You can increase the initial and maximum size of the Central/RAS heap by configuring the **central-wrapper.conf** and **ras-wrapper.conf** files (located under **<installation_folder>/<central or ras>/conf/**).

```
wrapper.java.initmemory=<value in MB>
```

```
wrapper.java.maxmemory=<value in MB>
```

We recommend configuring Central to 4GB and RAS to 2GB.

- **Worker threads** - By default, each node has 20 worker threads. If your flows have a large number of parallel or multi-instance lanes, we recommend gradually increasing this number, over several iterations.

You can configure the number of worker threads in the **central-wrapper.conf** and **ras-wrapper.conf** files (located under **<installation_folder>/<central or ras>/conf/**).

Open the file and edit the `-Dcloudslang.worker.numberOfExecutionThreads` property.

This is a new property, which was not supported in versions earlier than 10.22. Therefore, it should be added, if it doesn't exist, as follows:

```
wrapper.java.additional.<next available number>=
```

```
-Dcloudslang.worker.numberOfExecutionThreads=<new value>
```

- **Database connections** - You can increase the number of database connections in the **database.properties** file (located under `<installation_folder>/central/conf/`).

Edit the `db.pool.maxPoolSize` property. We recommend changing it to 100 connections.

For more information about system tuning, see the *Tuning Guide*

Database Maintenance

Database maintenance is important in order to maintain the efficiency and throughput of OO.

For detailed information on database housekeeping (index rebuilding, table statistics updated, and so on), see the OO *Database Guide*. "[Set Up Database Environment](#)" below

Set Up Database Environment

This documents shows the database administrator how to configure four different types of database, to be used with OO.

Introduction to Preparing the Database Environment

This section contains information about the types of databases used with (OO).

Overview

The term “database” may be interpreted in several ways, depending on the database vendor/technology used. In Oracle, the term “database” relates to a collection of files containing data and metadata. A single Oracle database may contain one or more schemas (and users). A Microsoft SQL Server “database” is closer in definition to Oracle’s “schema” than to Oracle’s “database”.

In order to avoid confusion, this document will use the following terms:

- Instance/server – the software and memory structures providing RDBMS services
- Database – the entity containing tables, views, indexes, and so on

OO requires a single database to be created. This database may co-exist with other databases contained in a database server.

You can set up an OO database on one of the following database server types:

- Microsoft SQL Server Standard/Enterprise (2008 R2/2012/2014/2016)
- Oracle 11gR2 Standard/Enterprise Server, including RAC environment
- Oracle 12cR1 Standard/Enterprise Server - regular instance (non-CDB), including RAC environment
- Postgres (9.1/9.2/9.3/9.4/9.5/9.6)
- MySQL Community/Standard/Enterprise Server (5.5/5.6/5.7)

See the appropriate deployment chapter below for details:

- ["Deploying and Maintaining a Microsoft SQL Server Database" on page 33](#)
- ["Deploying and Maintaining an Oracle Database" on page 46](#)
- ["Deploying and Maintaining a MySQL Database" on page 57](#)
- ["Deploying and Maintaining a Postgres Database" on page 66](#)

Appendices contain additional information that is pertinent to all database types.

Language Support

OO can be installed and used in any supported language (English, French, German, Spanish, Japanese, and Simplified Chinese). Databases and database servers should be properly configured in order to support the desired language.

If you are using OO in a multilingual environment, it is preferable that you configure your database to use the Unicode character set. See the relevant deployment chapter for detailed instructions.

If you have user input in two languages apart from English (for example, German and Chinese) then MS SQL should not be used. You should use an alternative database such as Oracle, MySQL, or Postgres with the recommended Unicode configuration for OO.

Important Notes

- This document is intended for trained database administrators. If you are not familiar with the type of database you wish to use, or you feel that you lack the relevant knowledge required in order to create and configure an OO database, see the database vendor's documentation and make sure you fully understand each action you perform as described in this section.

- OO database connectivity relies on Java JDBC. If your environment requires specific adaptations or security measures, see the JDBC documentation (or database vendor documentation) to find out exactly how the JDBC connection URL should be formatted.
- This document describes the required database settings for OO. Settings that are not specified in this document can be left with default values or configured by your organization's DBA.

Using Database Clusters

Database clusters help to make the OO system more robust, by protecting the database from several types of failures.

OO does not provide database cluster-related features of its own, such as database connection failover. It relies on the functionality of the JDBC package being used, and the database cluster environment capabilities, such as SCAN/AG listeners.

You can install OO in conjunction with any type of database cluster environment that satisfies the following conditions:

- Connection pooling is supported
- A single, valid database URL can be provided
- It provides reliable transaction handling during failover (a single, complete transaction must either completely fail or succeed)

The OO installer does not provide database cluster-specific installation options. During installation, you need to provide the installer a simplistic / "regular" form of database connection. Database URL adaptations usually have to be made after the system is already installed in the **database.properties** configuration file.

Database Security

The OO database is at the heart of the OO system, containing the OO system configuration and possibly corporate-sensitive data. It is highly recommended that you enforce strict security settings on your database by following these guidelines.

Harden your database server by following the recommendations from the database vendor and the operating system vendor. Server hardening may include, among other items:

- Restrict SSH access on Linux servers to a well-controlled set of OS users
- Enforce a strict OS user password policy – password length, complexity, lockout-policy, and so on
- Set up an audit system to detect and report break-in attempts

Harden your database user/login accounts:

- Enforce the password policy – password length, complexity, lockout-policy, and so on
- Restrict access to administration-level accounts
- Set up an audit system to detect and report break-in attempts

Database encryption at file-system level is supported as long as it does not compromise OO's performance, and completely transparent to OO. For example, Oracle TDE may be used as long as it does not cause any performance degradation or prevent LOB free-space reclaim.

Database Schema

OO's databases are always created using the default database schema. For SQL Server, OO uses the **dbo** schema, and for PostgreSQL - the **public** schema. Currently, no other schemas are supported by OO.

OO Database Sizing

OO 10.x keeps all flow and step execution data in the database by default, enabling comprehensive debugging of previous flow runs. As a result, the database size grows in accordance to the system throughput and flow complexity.

In version 10.22 and later, the amount of data kept in the database can also be controlled by setting the system persistency level. See the *OO Tuning Guide* for more details.

This section will help you prepare for installing OO. By estimating your system scale (small/standard/enterprise), you will be able to derive the amount of disk space required, derive the amount of memory (RAM) used by the database, and determine additional database installation parameters.

Step 1: Estimate the system scale according to complexity:

System Criteria\Scale	Small	Standard	Enterprise
Average number of steps per	50 or less	100 or more	1000 or more

System Criteria\Scale	Small	Standard	Enterprise
flow			
Average flow duration	Seconds to minutes	1 hour or more	Up to 24 hours
Average payload size per flow ¹	1 KB or less	1 MB or more	4 MB or more

[1] – The term "payload" refers to flow/step data of considerable size. For example, a CSV or XML file used as flow input, a large JSON object / REST API call data.

Step 2: Estimate the system scale according to concurrency/frequency:

System Criteria\Scale	Small	Standard	Enterprise
Average number of flows per day	Less than 100	300 or more	1000 or more

The following table provides disk space and memory requirements for different OO deployment scales:

System Scale\Parameter	OO Database Disk Size Requirements	Memory
Small	50 GB	4 GB
Standard	500 GB	8 GB
Enterprise	Between 500 GB and 2 TB	12 GB

Notes:

- Disk space and memory values are estimates. Actual disk space and memory consumption vary, depending on the database vendor and database server configuration.
- Memory (RAM) reflects recommended database memory, not the overall amount of memory available on the database machine (virtual or physical server).
- Disk space reflects the amount of disk space required for day-to-day operation of the OO system and reasonable historical data retention—not including long-term database backups.

It is highly recommended to regularly purge historical data using the stored procedure published on ITOM Marketplace and the database purge flows provided in the HPE Solutions content pack, in order to keep the database size within your environment limitations.

- The amount of additional disk space required for keeping OO database backups depends on the backup policy (frequency and retention period).

Hardware Requirements

The following table describes the hardware (CPU and memory) requirements recommended for the each of the database servers.

Note: The memory values reflect database memory consumption as part of the entire machine's memory.

Database\Scale	Small/Standard				Enterprise			
	CPUs		RAM		CPUs		RAM	
	Small	Standard	Min	Rec	Min	Rec	Min	Rec
SQL Server	2	4	4 GB	8 GB	4	12	8 GB	12 GB
Oracle	2	4	4 GB	8 GB	4	12	8 GB	12 GB
MySQL	2	4	4 GB	8 GB	4	12	8 GB	12 GB
Postgres	2	4	4 GB	8 GB	4	12	8 GB	12 GB

Min = Minimal value, Rec = Recommended value

In addition to the generalized hardware requirements above, refer to the appropriate hardware requirements and software requirements sections per database.

Deploying and Maintaining a Microsoft SQL Server Database

In order to deploy OO using Microsoft SQL Server, you must have an existing SQL Server database service. If you need to create a new database service, see the relevant documentation provided by Microsoft, because this information is not included within this section. However, this section contains recommendations for the SQL Server configuration.

It is also recommended to use the SQL Server Agent service in order to comfortably schedule data purging and index maintenance jobs.

This chapter includes the following sections:

Workflow for Microsoft SQL Server Deployment

To deploy OO using Microsoft SQL Server, perform the following steps:

1. **Review sizing guidelines.** For details, see "OO Database Sizing" in ["Introduction to Preparing the Database Environment"](#) on page 28.
2. **Review Hardware and Software Requirements.** For details, see ["System Requirements for Microsoft SQL Server"](#) below.
3. **Configure Microsoft SQL Server.** For details, see ["Configuring SQL Server"](#) on page 37.
4. **Create an OO database on Microsoft SQL Server.** For details, see ["Manually Creating an OO Database on Microsoft SQL Server"](#) on page 38.
5. (Optional) **Set up Windows authentication.** For details, see "Using Windows Authentication to Access Microsoft SQL Server Databases" in ["Additional Guidelines for Microsoft SQL Server"](#) on page 76. This step is only relevant if you are using Windows authentication instead of SQL Server authentication.

System Requirements for Microsoft SQL Server

This section describes the system requirements for working with Microsoft SQL Server in conjunction with OO.

Hardware Requirements

For OO database sizing guidelines and hardware requirements, see "OO Database Sizing" and "Hardware Requirements" in ["Introduction to Preparing the Database Environment"](#) on page 28.

For Microsoft SQL Server hardware requirements, see the relevant installation section for your Microsoft SQL Server release and operating system.

Software Requirements

The following table lists the Microsoft SQL Server releases supported by OO:

Microsoft SQL Server Database Releases			
Version	Type	32/64Bit	Service Pack
2016	Standard	64 Bit	1
	Enterprise	64 Bit	1

Microsoft SQL Server Database Releases			
Version	Type	32/64Bit	Service Pack
2014	Standard	64 Bit	1
	Enterprise	64 Bit	1
2012	Standard	64 Bit	2
	Enterprise	64 Bit	2
2008 R2	Standard	x64	3
		x86	3
	Enterprise	x64	3
		x86	3

Only the listed service packs should be installed. Newer service packs are also supported unless stated otherwise in the *OO Release Notes*.

See the Microsoft documentation for supported platforms.

Examples of Tested Deployments

The following table lists the deployment environments that have been rigorously tested by quality assurance personnel.

Database Release			Operating System
Version	32/64Bit	Service Pack	
Microsoft SQL Server 2014 Enterprise Edition	64 Bit	1	Windows 2012 Standard Edition 64 Bit
Microsoft SQL Server 2012 Enterprise Edition	64 Bit	2	Windows 2012 Standard Edition 64 Bit
Microsoft SQL Server 2008 R2 Enterprise Edition	64 Bit	3	Windows 2012 Standard Edition 64 Bit

Language Support

In Microsoft SQL Server, unlike other databases, OO database does not use Unicode collation.

Use one of the following collations depending on your OO installation language:

Language	Database Collation
English	SQL_Latin1_General_CP1_CS_AS
Japanese	Japanese_Unicode_CS_AS
Simplified Chinese	Chinese_Simplified_Stroke_Order_100_CS_AS
German	SQL_Latin1_General_CP1_CS_AS
French	French_100_CS_AS
Spanish	SQL_Latin1_General_CP1_CS_AS

Note: When an MS SQL database is created with a case-sensitive collation, object names such as tables, keys, and so on become case-sensitive as well.

For example, for the OO_STEP_LOG_BINDINGS table, a command such as `SELECT * FROM oo_step_log_bindings` is seen as using an invalid object name.

If you are currently using a different collation, it is highly recommended that you change your OO database collation to one of the collations above, in order to support future OO versions.

The following procedure is an example of how to change an existing database collation:

1. Connect to your database server using an administrative login (for example, “sa”).
2. Disconnect all existing sessions to the specific OO database that you intend to change.

Important! The command will fail unless this database has exactly 0 sessions/connections.

3. Run the following code (change `my_database` to the actual name of the database):

```
USE [master]
GO
ALTER DATABASE [my_database] COLLATE Japanese_Unicode_CS_AS
GO
```

Note: This change does not alter existing column collations. Any new column or table will use the new collation by default from this point on. The new collation sorting rules are applied immediately. In other words, the new collation affects sorting behavior and future data, rather than existing data.

If you have user input in two languages apart from English (for example, German and Chinese) then MS SQL should not be used. You should use an alternative database such as Oracle, MySQL, or Postgres with Unicode configuration.

Configuring SQL Server

This section contains information on Microsoft SQL Server and database configuration settings.

You can install an OO database in any SQL Server environment including clustered environments.

Legend:

- **Mandatory** configuration options/values appear in **bold/orange** font.
- **Recommended** configuration options/values appear in **bold/purple** font.
- Supported configuration options/values appear in normal font, and may show as a comma-separated list.
- *Comments* appear in *italic* font.

Microsoft SQL Server 2008R2, 2012, 2014, and 2016	
Server Options/Features	
Configuration Item	Supported Configuration Options
Server Configuration Options	Defaults, unless instructed otherwise
Instances	Default, Single
Authentication Mode	Mixed, Windows ¹
Full-Text Search	(Not required for OO)

[1] At the moment, the OO 10.x installer supports only SQL authentication. Windows authentication can be configured later.

Microsoft SQL Server 2008R2, 2012, 2014, and 2016			
Instance/Server Options			
	Mandatory	Recommended	Supported
Server Collation		SQL_Latin1_General_CP1_CS_AS	Any Collation
Network Libraries	Server: TCP/IP Client: TCP/IP		
Concurrent Connections	>= 800	0 (<i>unlimited</i>)	
Max Server Memory	> 4GB	2,147,483,647 (<i>default, unlimited</i>) <i>Allocate 4-12 GB depending on system scale according to the sizing guide</i>	

Microsoft SQL Server 2008R2, 2012, 2014 and 2016			
Database Options			
	Mandatory	Recommended	Supported
Collation	Any collation listed in the “Language Support” section in “System Requirements for Microsoft SQL Server” on page 34.		
Recovery Model		Full	Simple, Full
Allow Snapshot Isolation	True		
Is Read Committed Snapshot On	True		
Auto Shrink	False		
Auto Create Statistics	True		

Note: SQL Server uses a READ-COMMITTED transaction isolation without snapshot isolation by default. It is mandatory to set the **Allow Snapshot Isolation** and **Is Read Committed Snapshot On** flags. OO fails to operate properly using any other type of transaction isolation.

Manually Creating an OO Database on Microsoft SQL Server

During OO setup, a new database can be created automatically by the OO installer or an existing database can be used.

If during installation, you are authorized to connect to the database server as **sysadmin** (that is, connect as “sa”), use the installer’s “Create the database/schema” option, and you can skip this section.

This section describes the procedure for manually creating an OO database, login, and user on Microsoft SQL Server.

Note: Only the database, login and user are created at this point; objects such as tables and

indexes are not created yet. These objects are created when OO starts for the first time.

This section is relevant for you if, for example, due to security restrictions, you do not wish to use login/user credentials with elevated privileges during the OO installation. In such a case, you (or your organization's DBA) should create the database, login, and user first, and then let the OO installer connect to the pre-existing database using "lower" privileges.

To create a database, you must connect to the SQL Server instance using a login that has **CREATE DATABASE** permission.

- Members of the sysadmin server roles automatically have **CREATE DATABASE** permission, and are also mapped to dbo in all databases.
- Perform the following procedures only if you are an experienced Microsoft SQL Server database administrator.
- If you prefer to use the database creation wizard/GUI, make sure you select all options that correspond with the T-SQL code presented below. For example, make sure you set **Allow Snapshot Isolation** to TRUE under the **Options** page/**Other Options** pane/**Miscellaneous** tab in the New Database dialog box.
- Not all database creation options are specified—only those that differ from the default value. When in doubt, use default values.

Selecting a specific collation also means that an OO system that uses SQL Server is limited to the set of languages supported by that specific collation. For example, if the **SQL_Latin1_General_CP1_CS_AS** collation is used, English, German, and Spanish characters may be used, but Japanese characters may not. If **Japanese_Unicode_CS_AS** is used, French accent characters will not be presented properly. For the complete specification of each collation, see the Microsoft SQL Server documentation.

To create a database:

1. Log in to Microsoft SQL Server as "sa" or any other login with a **sysadmin** role or the **CREATE DATABASE** permission.
2. Run **<installation folder>/central/bin/sql/MS-SQL/mssql_create_oo_db.sql** and verify that the database, login, and user were created successfully. Adapt the script parameters to match your environment.
3. (Optional) In order to verify that database objects can be created by the new login and user, connect to the database server using OO's newly created login role credentials and run **<installation folder>/central/bin/sql/MS-SQL/mssql_create_test_table.sql**.

Verify that the script ran correctly and that no errors are shown.

Manually Creating Database Objects

Once the database, login, and user are in place, the database objects (tables, indexes, and so on) are created when the OO service starts and connects to the database for the first time.

You could create the database objects manually if the user created for OO does not have permissions to create or modify database objects (limited to DML operation only).

It is highly recommended to grant the OO database user DDL-related privileges, in order to allow OO to perform upgrade operations without manual intervention, as these operations sometimes require modifications to the database structure.

To manually create the database objects:

1. Extract the **mssql.sql** files from the OO installation zip file under **docs\sql**.
2. Log in to Microsoft SQL Server as the relevant user, who is permitted to create and modify database objects in the OO database.
3. Run the **mssql.sql** file and verify that no errors occurred.

Microsoft SQL Server Database Maintenance

This section describes the various maintenance tasks that are recommended for OO databases created on Microsoft SQL server, such as backing up the database, checking database integrity, handling index fragmentation, and monitoring the database.

Backing up the OO Database

Microsoft SQL Server databases are either configured for the **Full** recovery model, or the **Simple** recovery model. You can back up an OO database using either one of these models. As OO keeps all configuration and operational history in a single database, always backup the complete database.

Consider the following guidelines when you create your backup plan for OO:

Backup method:

The backup method depends mainly on business considerations—how much information “may” be lost? What is the maximum time for system recovery? If you need to be able to perform point-in-time recovery, and only “allow” a few hours-worth of data loss, use the full recovery model and perform full or differential backups daily, and transaction log backup every N hours depending on your business requirements.

If your organization is more tolerant to data loss, you can use the simple recovery model and perform a daily or weekly full backup.

Backup frequency:

Daily backup is recommended, especially if you are using/modifying OO on a daily basis.

You should back up once a month at the very least.

Timing:

Schedule backup for the time when OO is least active.

Retention:

Retention depends on your business guidelines and regulations.

Creating a Maintenance Plan

Maintaining an OO database includes rebuilding the index and reclaiming free space. Use the scripts and tools described in this section, in order to keep the OO database in good shape.

Supplied Utilities for OO Database Housekeeping

OO provides a set of scripts for index maintenance, statistics maintenance, and history purging. These scripts create stored procedures that you can tune and schedule to run periodically.

OO provides a set of scripts for index maintenance, statistics maintenance, and history purging. These scripts create stored procedures that you can tune and schedule to run periodically.

Note: SQL scripts are available in the Central server installation under the **central/bin/sql** folder.

It is recommended to use these procedures, but you can also use other methods in accordance with company policy, as long as indexes and statistics are kept in good shape.

Note that index rebuilding online (without OO system downtime) requires an enterprise grade database. Make sure you are running an enterprise version of Microsoft SQL Server before attempting online index rebuilding.

Also, note that maintenance activity usually consumes additional resources from the database server. This is why it is important to schedule maintenance activity at the times when OO is least active.

Utility for Maintaining Indexes and Statistics

Download the latest **HPE_OO_DB_Maintenance.zip** pack from ITOM Marketplace and unpack it to a temporary folder.

To install and use the OO maintenance stored procedures:

1. Log in to Microsoft SQL Server as “sa” or any member of the **sysadmin** role and execute: `<installation folder>/central/bin/sql/MS-SQL/mssql_grant_view_server_state.sql` in order to give the OO user the ability to access **dm_os_performance_counters** dynamic management view (DMV).
2. Edit the **OOIndexMaintenance.sql** and (optionally) **OO_DB_MAINTENANCE_LOG.sql** scripts under `<temporary folder>/HPE_OO_DB_Maintenance/MS-SQL/10.60_to_10.80/` and replace “USE <your_db_name_here>” at the beginning of the script with your actual OO database name. For example, if your database name is “OOPROD”, replace it with “USE OOPROD”.

Do not skip this step; otherwise, the set of procedures will not be created in the correct database.
3. Log in to Microsoft SQL Server as the OO user.
4. Run the **OOIndexMaintenance.sql** and (optionally) **OO_DB_MAINTENANCE_LOG.sql** T-SQL scripts in the specified order and verify that the new objects were created successfully.
5. Tune your stored procedure according to the explanations below.

`<installation folder>/central/bin/sql/MS-SQL/mssql_execute_index_maintenance.sql` provides an example of how the index maintenance procedure can be used.

- **@FragmentationXXX** sets the script’s fragmentation level sensitivity and course of action in each case. These threshold levels and subsequent actions are recommended by Microsoft’s documentation. Tune these values with caution.
- **@SortInTempdb** (once set to “Y”) lets you perform sorting operations during index reorganization/rebuild in tempdb rather than in memory, for better performance. If you choose to use this option, make sure you have sufficient free space in **tempdb**.
- **@Indexes** is a filter for including/excluding indexes in the maintenance operation. It is recommended to keep this filter as is, to analyze all indexes.
- **@TimeLimit** is the timeout in seconds for the maintenance operation to complete. Set it in accordance with your maintenance window boundaries, if applicable.
- **@LockTimeout** is the timeout in seconds to wait for object lock. Once expired, the specific operation fails and the procedure continues to the next object.
- **@LogToTable** determines whether maintenance operation results should be logged to a table. This lets you keep track of the maintenance operations and helps procedure debugging.

Note that it requires creating the table using the **OO_DB_MAINTENANCE_LOG.sql** script.
- **@Execute** determines whether actual operations (such as index rebuild) are performed or not. If this parameter is set to ‘N’, the procedure performs a “dry run” and shows an analysis of the relevant objects.

Utility for Purging Historical Data

To install and use the OO history purging stored procedure:

1. Edit the **OOPurgeHistory.sql** script under **<temporary folder>/HPE_OO_DB_Maintenance/MS-SQL/10.60_to_10.80/** and replace "USE <your_db_name_here>" at the beginning of the script with your actual OO database name. For example, if your database name is "OOPROD", replace it with "USE OOPROD".

Do not skip this step; otherwise, the procedure will not be created in the correct database.

2. Log in to Microsoft SQL Server as the OO user.
3. Run the **OOPurgeHistory.sql** T-SQL script and verify that the new objects were created successfully.
4. Tune your stored procedures according to the explanations below.

<installation folder>/central/bin/sql/MS-SQL/mssql_execute_purge_stored_procedure.sql provides an example of how the index maintenance procedure can be used.

- **@PurgeExecutionsOlderThan** determines how many days are kept, (protected) relative to the time the procedure starts running. By default, 90 days are kept. Older data is deleted, starting with the oldest records.
- **@PurgeExecutionsInBatchesOf** determines how many flows are handled together. Smaller values imply smaller, more frequent transactions, and higher values imply less frequent, larger transactions. 1000 is recommended for most systems.
- **@ShouldPurgeExecutionSummary** determines if the **OO_EXECUTION_SUMMARY** table should be purged. The default value is "0" (do not purge this table). It is recommended to keep data in this table as it does not consume a lot of space. Use "1" only if you want to completely remove any reference to the relevant flows.
- **@verbose** determines the verbosity level. "0" corresponds to "quiet" output, "1" corresponds to normal output, and "2" prints out detailed information.
- **@StopPurgingProcessAfter** is the timeout in hours for the operation to complete. Set it in accordance with your maintenance window boundaries if applicable.
- **@DeepClean** determines whether deep cleansing is performed. For example, searching for "orphan" records that may bloat the database unnecessarily. The default is "0" (off). Note that setting this flag to "1" prolongs the procedure run time, but the timeout limit is still imposed.
- **@DisableIndexes** determines whether certain database indexes are disabled for the duration of the purge operation. These indexes are rebuilt at the end of the purge operation. Using this feature speeds up the purge on the expense of flow drill-down (step-level debugging).

Upgrade/Rollback Instructions

Prior to any upgrade or rollback of OO, take the following steps:

1. Purge non-essential data from the database.

Upgrade and rollback operations convert flow execution data. Less data means a shorter upgrade/rollback. If all flow execution data is essential, perform a complete database backup first and then purge.
2. Back up your database, preferably after the database was purged.
3. Schedule the operation to the database server's quiet period.
4. Verify that no shared resources are exhausted – temporary space, transaction log, and similar resources.

Following any upgrade or rollback of OO, take the following steps:

1. Install relevant versions of maintenance and purge stored procedures.

Stored procedures and purge flows are not upgraded/rolled back by the OO upgrade utility.
2. Reschedule database maintenance and purge routines (stored procedures and/or purge flows).

Database Server Upgrade from 2012 or 2008R2 to Server 2014 or Server 2016

If the OO 10.60x or higher database server is being upgraded from 2012 or 2008R2 to 2014, you must modify the OO Central server configuration in order to use a different JDBC connector jar.

SQL Server 2008R2 and 2012 are interfaced by OO using JTDS 1.3.0 connector, while 2014 is interfaced using Microsoft JDBC 4.2.

If your SQL Server database server is being upgraded to 2014:

1. Edit central server **database.properties** file located under the **central\conf** folder.
2. Modify the **db.driver** value as follows:

From:

net.sourceforge.jtds.jdbc.Driver

To:

```
com.microsoft.sqlserver.jdbc.SQLServerDriver
```

3. Modify the **db.url** value as follows:

From:

```
jdbc:jtds:sqlserver://db_host:db_port/db_name;sendStringParametersAsUnicode=true
```

To:

```
jdbc:sqlserver://db_host:db_port;databaseName=db_name;sendStringParametersAsUnicode=true
```

If you have edited your JDBC URL and used additional options, use Microsoft's JDBC 4.2 document in order to format the new URL correctly.

4. Repeat all above steps for each central server in the OO system.
5. Restart the OO Central servers.

If OO is being upgraded to 10.6x or higher and the SQL Server database is being upgraded to 2014 simultaneously, follow these steps:

1. Complete the OO upgrade procedure using the current database (2008R2 or 2012).
Follow all the instructions in the sections above as usual – purge beforehand, backup, etc. It is highly recommended to perform a sanity test for the OO system using the current database.
2. Make a backup of the OO database (full, copy-only backup is preferred). Backup OO's login role, if necessary.
3. Upgrade the SQL Server database server to 2014.
4. Restore OO's database from the backup if necessary (optional).

If OO's database was upgraded along with the database server, there is no need to restore it. Reconfigure the login role authentication and database ownership if necessary.

5. Perform the steps as described in "[Microsoft SQL Server Database Maintenance](#)" on page 40

At this point, you must configure Central servers to use Microsoft JDBC instead of JTDS.

Always On Support

Microsoft SQL Server 2008R2 and 2012 "Always On" is supported by OO as a legacy client^[1] to implement a high availability and/or disaster recovery solution.

[1] OO does not support the “MultiSubnetFailover” feature for SQL Server 2008R2 and 2012, and requires that **RegisterAllProvidersIP** is set to 0 in the cluster/listener configuration.

Connection to SQL Server 2014 “Always On” is supported including the “MultiSubnetFailover” feature, but has not been certified for OO.

See [“Additional Guidelines for Microsoft SQL Server” on page 76](#) for additional information.

High Availability Setup

As OO requires a single point of contact, it is essential that an Availability Group Listener (AG listener) exist in the high availability setup. OO’s database connection is defined by a single URL - containing a single hostname.

Normally, OO interfaces with the AG listener in order to connect to the availability group primary replica (a READ-WRITE capable instance).

In the case of a database failover scenario, while database connectivity is lost, OO will repeatedly attempt to connect to the database until a new READ-WRITE capable instance accepts the connection. The redirection of the configured database hostname to a different IP address should be done by the AG Listener and cluster environment.

Note: OO has no use for a read-only database (thus, has no use for read-only routing).

Disaster Recovery Solution

In a disaster recovery setup, OO should be halted, reconfigured (edit the **database.properties** file), and then started again once the new database instance is available.

See the *OO Installation, Upgrade, and Configuration guide* for details on reconfiguring the database connection details.

Deploying and Maintaining an Oracle Database

In order to deploy Operations Orchestration using Oracle, you must have an existing Oracle database service. If you need to create a new database instance/service, see the relevant documentation provided by Oracle, because this information is not included within this section. However, this section contains recommendations for the Oracle instance configuration.

This chapter includes the following sections:

Workflow for Oracle Deployment

To deploy OO using Oracle, perform the following steps:

1. **Review sizing guidelines.** For details, see "OO Database Sizing" in ["Introduction to Preparing the Database Environment" on page 28.](#)
2. **Review Hardware and Software Requirements.** For details, see ["System Requirements for Oracle" below.](#)
3. **Configure an Oracle Database.** For details, see ["Configuring an Oracle Database" on page 49.](#)
4. **Create a Database.** For details, see ["Manually Creating an OO Database on an Oracle Instance" on page 51.](#)
5. (Optional) **Connect OO to an Oracle RAC environment.** This step is only relevant if you are using OO in an Oracle RAC environment.

System Requirements for Oracle

This section describes the system requirements for working with Oracle in conjunction with OO.

Hardware Requirements

For OO database sizing guidelines and hardware requirements, see "OO Database Sizing" and "Hardware Requirements" in ["Introduction to Preparing the Database Environment" on page 28.](#)

For Oracle hardware requirements, see the relevant installation section for your Oracle release and operating system.

Software Requirements

The following table lists the Oracle releases supported by OO:

Oracle Releases			
Version	Type	32/64Bit	Patch Set
11g R2	Standard	64 Bit	11.2.0.1 – 11.2.0.4 ^[1]
	Enterprise	64 Bit	11.2.0.1 – 11.2.0.4 ^[1]
12c R1 Regular instance, non-CDB	Standard	64 Bit	12.1.0.1 - 12.1.0.2
	Enterprise	64 Bit	12.1.0.1 - 12.1.0.2

Only the listed patch sets should be installed. Newer patch sets are also supported unless stated otherwise in the OO *Release Notes*.

[1] It is highly recommended to apply Oracle's "DATABASE PATCH SET UPDATE 11.2.0.4.6", which solves an issue with rollback to previous OO versions.

See the Oracle documentation for supported platforms.

Note that Oracle 12c R1 is only supported in its regular, backward-compatible instance form. OO 10.x does not support connecting to an Oracle 12c container database (CDB).

The Oracle 12c RAC environment is supported in backward-compatible form. New features such as multiple cluster subnets for SCAN listener are not supported.

OO is compatible with Oracle Data Guard protected databases. For OO to work with a database protected by Oracle Data Guard, ensure that the Data Guard protection mode is set to "Maximum protection". The Data Guard environment must be configured for Fast-Start Failover to limit the down time of the OO service during failovers. The JDBC URL requirements are similar as those for Oracle RAC.

Note: During failover, OO may function in read only mode (that is, no new flow executions can occur) for a short period of time while the database change occurs. Also, during failover, historical data may be out of sync. This will be resolved when the failover completes.

Oracle Connector

The jdbc connector .jar file is not provided with the OO installation. You must download it and place it under each Central <installation folder>/central/lib folder.

Verify that the connector version you are using is fully compatible with your database server version. The recommended driver is the Oracle JDBC driver version 7-12.1.0.2

Examples of Tested Deployments

The following table lists the deployment environments that have been rigorously tested by HPE quality assurance personnel.

Database Release			Operating System
Version	32/64Bit	Patch Set	
Oracle 11g R2 Enterprise Edition	64 Bit	11.2.0.4.6	Red Hat Enterprise Linux 6.5 64 Bit

Database Release			Operating System
Version	32/64Bit	Patch Set	
Oracle 11g R2 Enterprise Edition	64 Bit	11.2.0.4.0	Windows 2012 Standard Edition 64 Bit
Oracle 12c R1 Enterprise Edition	64 Bit	12.1.0.2	Red Hat Enterprise Linux 6.5 64 Bit

Language Support

The Oracle instance character set should be set to **AL32UTF8**. This enables using any Unicode character (and practically all common characters in all languages).

Configuring an Oracle Database

This section contains information on the Oracle instance and database configuration settings.

You can install an OO database in an Oracle clustered environment (Oracle RAC or other).

Legend:

- **Mandatory** configuration options/values appear in **bold/orange** font.
- **Recommended** configuration options/values appear in **bold/purple** font.
- Supported configuration options/values appear in normal font, and may show as a comma-separated list.
- *Comments* appear in *italic* font.

Oracle Database 11gR2 / 12cR1			
Instance/Server Options			
Instance configuration options	Defaults, unless instructed otherwise		
	Mandatory	Recommended	Supported
PROCESSES	>= 500		
SESSIONS	>= 555		
TIMED_STATISTICS		TRUE	TRUE, FALSE
OPEN_CURSORS	>= 900		

Oracle Database 11gR2 / 12cR1			
Instance/Server Options			
Instance configuration options	Defaults, unless instructed otherwise		
	Mandatory	Recommended	Supported
Shared/Dedicated Server		Dedicated	Dedicated, Shared
UNDO_MANAGEMENT		AUTO	Automatic, Manual
Undo size	>= 4GB	6GB – 10GB	
Memory Management		ASMM	AMM, ASMM
MEMORY_TARGET		0 (disabled)	>= 5G (for AMM)
SGA_TARGET		8G – 12G	>= 4G (for ASMM)
SGA_MAX_SIZE		8G – 12G	>= 4G (for ASMM)
PGA_AGGREGATE_TARGET		1G – 2G	>= 500M (for ASMM)

- Note that all values reflect resources required by OO. If OO shares the Oracle instance with other users, these values should be added to whatever is currently consumed by others.
- See the sizing guide to determine values displayed as range.

Oracle Database 11gR2 / 12cR1			
Instance/Server Options			
	Mandatory	Recommended	Supported
File system			ASM, Any
Storage options		Locally managed tablespace	
		Automatic segment space management (ASSM)	
		Automatic local extent management	
ARCHIVELOG mode		ARCHIVELOG	ARCHIVELOG, NOARCHIVELOG
Redo Log total size	>= 600MB	1GB	

- Note that all values reflect resources required by OO. If OO shares the Oracle instance with other users, these values should be added to whatever is currently consumed by others.

- See *theOO Sizing Guide* to determine values displayed as range.

Manually Creating an OO Database on an Oracle Instance

During OO setup, a new database can be created automatically by the OO installer or a pre-existing database can be used.

If during installation, you are authorized to connect to the database server as dba (connect as "SYSTEM"), use the "create the database/schema" option, and you can skip this section.

Note: In some cases, the term "database" is used but in the case of Oracle, it should be interpreted as "user".

This section describes the procedure for manually creating an OO database in an Oracle instance.

Note: Only the database is created at this point; objects such as tables and indexes are not created yet. These objects are created once OO starts for the first time.

This section is relevant for you if, for example, due to security restrictions, you do not wish to use user credentials with elevated privileges during the OO installation. In such a case, you (or your organization's DBA) should create the user (database) first, and then let the OO installer connect to the pre-existing database using basic privileges.

To create a database, you must connect to the Oracle instance using a login that has **CREATE USER** system privilege—for example, system user.

- Any user with the DBA role has sufficient privileges to create the new user.
- Perform the following procedures only if you are an experienced Oracle database administrator.
- If you prefer to use the database creation wizard/GUI, make sure you select all options that correspond with the SQL code presented below.
- Not all database creation options are specified—only those that differ from the default value. When in doubt, use default values.

To create a database:

1. Log in to the oracle as "system" or any other user with a DBA role.
2. Edit the `<installation folder>/central/bin/sql/Oracle/oracle_create_oo_db.sql` SQL script, run it, and verify that the database was created successfully.
3. (Optional) In order to verify that database objects can be created by the new user, connect to the

Oracle instance as OO user, edit and run: **<installation folder>/central/bin/sql/Oracle/oracle_create_test_table.sql**

Verify that the script ran successfully, and no errors or warnings were shown.

Manually Creating Database Objects

Once the database is in place, the database objects (tables, indexes, and so on) are created when the OO service starts and connects to the database for the first time.

You can create the database objects manually if the user created for OO does not have permissions to create or modify database objects (limited to DML operation only).

It is highly recommended to grant the OO database user DDL-related privileges, in order to allow OO to perform upgrade operations without manual intervention, as these operations sometimes require modifications to database structure.

To manually create the database objects:

1. Extract the **oracle.sql** file from the OO installation zip file under **docs\sql**.
2. Edit the **oracle.sql** file and prefix each object with the OO user, to make sure the objects are created under the OO user.
3. Connect to the OO database as the relevant user, who is permitted to create and modify database objects in the OO database.
4. Run the **oracle.sql** file and verify that no errors occurred, and that all objects are created and owned by the OO user.

Connecting to Oracle Using SID or Service Name

Connecting to an Oracle database server requires you to specify either the SID (system ID) or Service Name. In version 10.20 and later, OO supports specifying the service name during installation (using either the Installation Wizard or a silent installation).

The following examples show how to format JDBC URL for each option, as it should appear in the Central **database.properties** file under the **central\conf** folder.

Connecting to an Oracle instance using the SID:

```
jdbc.url=jdbc\:oracle\:thin\:@DB_HOSTNAME_OR_IP\:PORT\:SID
```

Connecting to an Oracle instance using the Service Name:

```
jdbc.url=jdbc\:oracle\:thin\:@//DB_HOSTNAME_OR_IP\:PORT/SERVICE_NAME
```

Oracle Database Maintenance

This section describes the various maintenance tasks that are recommended for OO databases created on Oracle, such as backing up the database, checking database integrity, handling index fragmentation, and monitoring the database.

This section includes:

Backing up the Operations Orchestration Database

Oracle databases can be backed up using several tools, such as **expdp** and **RMAN**. An OO database can be backed up using any type of method/tool as long as the complete database is backed up.

Consider the following guidelines when you create your backup plan for OO:

Backup method:

The backup method depends mainly on business considerations—how much information “may” be lost? What is the maximum time for system recovery? If you need to be able to perform point-in-time recovery, and only “allow” a few hours-worth of data loss, use the full recovery model and perform full or differential backups daily, and transaction log backup every N hours depending on your business requirements.

Backup frequency:

Daily backup is recommended, especially if you are using/modifying OO on a daily basis.

You should back up once a month at the very least.

Timing:

Schedule backup for the time when OO is least active.

Retention:

Retention depends on your business guidelines and regulations.

Creating a Maintenance Plan

Maintaining an OO database includes rebuilding the index and reclaiming free space. Use the scripts and tools described in this section, in order to keep the OO database in good shape.

Supplied Utilities for OO Database Housekeeping

OO provides a set of scripts for index maintenance, statistics maintenance, and history purging. These scripts create a package containing stored procedures that you can tune and schedule to run periodically.

It is recommended to use these procedures, but you can also use other methods in accordance with company policy, as long as indexes and statistics are kept in good shape.

Note that index rebuilding online (without OO system downtime) requires an enterprise grade database. Make sure you are running an enterprise version of Oracle before attempting online index rebuilding.

Also, note that maintenance activity usually consumes additional resources from the database. This is why it is important to schedule maintenance activity at the times when OO is least active.

Utility for Maintaining Indexes and Statistics

Download the latest **HPE_OO_DB_Maintenance.zip** pack from ITOM Marketplace and unpack it to a temporary folder.

To install and use the OO maintenance stored procedures:

1. Log in to Oracle as “system” or any other user with a DBA role, edit, and execute **<installation folder>/central/bin/sql/Oracle/oracle_privs_for_index_maintenance.sql**. These system privileges are required in order to verify that the stored procedure created in the following steps has the explicit (not role-based) privileges to execute the index analysis and rebuild.
2. Open the **<temporary folder>/HPE_OO_DB_Maintenance/Oracle/10.60_to_10.80/** folder, log in to Oracle database as the user created for OO, and run the **HPE_OO_DB_MAINT.sql** script. Verify that the new package and procedures were created successfully.

In the same folder, **oracle_execute_IndexMaintenance.sql** provides an example of how this procedure can be used.

- **pMaxHeight (IN)** - The minimal index height threshold for index rebuilding. The Oracle documentation recommends 3. Smaller values may result in unnecessary rebuilding operations.
- **pMaxLeafsDeleted (IN)** - The minimal deleted leaves threshold for index rebuilding. The Oracle documentation recommends 15. Smaller values may result in unnecessary rebuilding operations.
- **pRebuild (IN)** - Should indexes be rebuilt (1) or only perform a dry-run (0). A dry-run will show only recommendations for index rebuilding.
- **pReturnValue (OUT)** - The number of rebuilt indexes.

Note: ONLINE index rebuilding should only be performed when the enterprise edition is used. Otherwise, the index rebuilding operation may lock tables and indexes and may interfere with the operation of OO.

Utility for Purging Historical Data

To install and use the OO history purging stored procedure:

1. Open the <temporary folder>/HPE_OO_DB_Maintenance/Oracle/10.60_to_10.80 folder, log in to Oracle database as the user created for OO, and run **HPE_OO_DB_MAINT.sql** script. Verify that the new package and procedures were created successfully.

In the same folder, **oracle_execute_PurgeHistory.sql** provides an example of how this procedure can be used.

- **pPurgeExecutionsOlderThan** determines how many days are kept (protected) relative to the time the procedure starts running. Older data is deleted, starting with the oldest records. This parameter has no default value and must be specified.
- **pPurgeExecutionsInBatchesOf** determines the maximum number of flows to handle in each batch. The default value is 10000.
- **pShouldPurgeExecutionSummary** determines if the **OO_EXECUTION_SUMMARY** table should be purged. The default value is "0" (do not purge this table). It is recommended to keep data in this table as it does not consume a lot of space. Use "1" only if you want to completely remove any reference to the relevant flows.
- **pVerbose** determines verbosity level. "0" corresponds to "quiet" output, "1" corresponds to normal output, and "2" prints out detailed information. Set it in accordance with your maintenance window boundaries if applicable.
- **pStopPurgingProcessAfter** is the timeout in hours for the operation to complete.
- **pDeepClean** determines whether deep cleansing is performed. For example, searching for "orphan" records that may bloat the database unnecessarily. The default is "0" (off). Note that setting this flag to "1" prolongs the procedure run time, but the timeout limit is still imposed.
- **@DisableIndexes** determines whether certain database indexes are disabled for the duration of the purge operation. These indexes are rebuilt at the end of the purge operation. Using this feature speeds up the purge at the expense of flow drill-down (step-level debugging).

Utility for Reclaiming LOB Space

Oracle databases are slow when it comes to freeing space occupied by LOB segments. Depending on the configuration of the database server, LOB segment space may not be released immediately after purge operations. As LOB space may have a significant impact on the database size, you can add a post-purge operation to force freeing of LOB segments.

The general structure of the command is:

```
ALTER TABLE FOO MODIFY LOB ( LOB_COLUMN_NAME ) (shrink space);
```

The list of tables and columns involved may vary from one OO version to another.

To install and use the OO LOB space reclaim stored procedure:

Open the <temporary folder>/HPE_OO_DB_Maintenance/Oracle/10.60_to_10.80 folder, log in to Oracle database as the user created for OO, and run **HPE_OO_DB_MAINT.sql** script. Verify that the new package and procedures were created successfully.

In the same folder, **oracle_execute_ForceLobShrink.sql** provides an example of how this procedure can be used.

Notes:

Several conditions must be met in order for the shrink operation to succeed:

- TABLESPACE and database tables must not be compressed.
- ROW MOVEMENT must be enabled for all tables.
- TABLESPACE must be configured as AUTO segment space management.
- LOBs storage must not be SECUREFILE.

See the Oracle documentation for the complete set of conditions.

Shrinking LOB segments of considerable size (tenths or hundreds of Giga-Bytes) may take some time to complete and may consume a considerable amount of database resources—CPU and I/O. Therefore, it is recommended to execute this process when the database server is least active.

Upgrade/Rollback Instructions

Rolling back OO 10.5x and later to previous versions requires patching Oracle 11gR2 database servers to 11.2.0.4.6. Apply Oracle's "DATABASE PATCH SET UPDATE 11.2.0.4.6" prior to rollback operation.

Prior to any upgrade or rollback of Operations Orchestration, take the following steps:

1. Purge non-essential data from the database.

Upgrade and rollback operations convert flow execution data. Less data means a shorter upgrade/rollback. If all flow execution data is essential, perform a complete database backup first and then purge.

2. Back up your database, preferably after the database was purged.
3. Schedule the operation to the database server's quiet period.
4. Verify that no shared resources are exhausted – temporary space, UNDO/REDO, and similar resources.

Following any upgrade or rollback of OO, take the following steps:

1. Install relevant versions of maintenance and purge stored procedures.
Stored procedures and purge flows are not upgraded/rolled back by the OO upgrade utility.
2. Reschedule database maintenance and purge routines (stored procedures and/or purge flows).

Deploying and Maintaining a MySQL Database

In order to deploy Operations Orchestration using MySQL, you must have an existing MySQL database. If you need to create a new database service, see the relevant documentation provided by MySQL, because this information is not included within this section. However, this section contains recommendations for the MySQL configuration.

This chapter includes the following sections:

Workflow for MySQL Deployment

To deploy Operations Orchestration (OO) using MySQL, perform the following steps:

1. **Review sizing guidelines.** For details, see "Operations Orchestration Database Sizing" in ["Introduction to Preparing the Database Environment" on page 28.](#)
2. **Review Hardware and Software Requirements.** For details, see ["System Requirements for MySQL" below.](#)
3. **Configure MySQL .** For details, see ["Configuring MySQL" on page 60.](#)
4. **Create an OO database on MySQL .** For details, see ["Manually Creating an OO Database on MySQL" on page 62.](#)

System Requirements for MySQL

This section describes the system requirements for working with MySQL in conjunction with OO.

Hardware Requirements

For OO database sizing guidelines and hardware requirements, see "OO Database Sizing" and "Hardware Requirements" in ["Introduction to Preparing the Database Environment" on page 28](#).

For MySQL hardware requirements, see the relevant installation section for your MySQL release and operating system.

Software Requirements

The following table lists the MySQL releases supported by OO:

MySQL Database Releases		
Version	Type	32/64Bit
5.5	Community	x86 32-bit
		x86 64-bit
	Standard	x86 32-bit
		x86 64-bit
	Enterprise	x86 32-bit
		x86 64-bit
5.6	Community	x86 32-bit
		x86 64-bit
	Standard	x86 32-bit
		x86 64-bit
	Enterprise	x86 32-bit
		x86 64-bit
5.7	Community	x86 32-bit
		x86 64-bit
	Standard	x86 32-bit
		x86 64-bit
	Enterprise	x86 32-bit
		x86 64-bit

Note: For MySQL 5.6.20 and 5.6.21, the requirements for the `innodb_log_file_size` have increased significantly.

For MySQL 5.6.1 - 19, the recommendation is 256M, but for MySQL 5.6.20 - 21, the recommendation is 2GB.

See the MySQL documentation for supported platforms.

MySQL Connector

The MySQL connector jar is not provided with the OO installation. Download and place it under each Central `<installation folder>/central/lib` folder.

Verify that the connector version you are using is fully compatible with your database server version. Currently, the only supported version is 5.1.35.

Upgrade Note: Because prior to version 10.60, MySQL JDBC 5.1.21 was in use, make sure to download and use 5.1.35 during the upgrade process.

Examples of Tested Deployments

The following table lists the deployment environments that have been rigorously tested by HPE quality assurance personnel.

Database Release			Operating System
Version	32/64Bit	Patch	
MySQL Server 5.6.19 Community Edition	64 Bit		Windows 2012 Standard Edition 64 Bit
MySQL Server 5.6.12 Community Edition	64 Bit		Red Hat Enterprise Linux 6.3 64 Bit

Language Support

The MySQL Server character set should be set to `utf8`. This lets you use any Unicode character (and practically all common characters in all languages). Note that the OO database uses the `utf8_bin` collation.

Configuring MySQL

This section contains information on MySQL and database configuration settings.

Legend:

- **Mandatory** configuration options/values appear in **bold/orange** font.
- **Recommended** configuration options/values appear in **bold/purple** font.
- Supported configuration options/values appear in normal font, and may show as a comma-separated list.
- *Comments* appear in *italic* font.

MySQL 5.5, 5.6 and 5.7			
Instance/Server Options			
Server Configuration Options	Defaults, unless instructed otherwise		
[mysqld]	Mandatory	Recommended	Supported
character-set-server	utf8		
collation-server	utf8_bin		
transaction-isolation	READ-COMMITTED		
max-allowed-packet	250M		
max-connections	>= 1000		
default-storage-engine	INNODB		
innodb_log_file_size (for MySQL 5.6.1 - 19)	256M		
innodb_log_file_size (for MySQL 5.6.20 - 21)	2GB		
innodb_log_file_size (for MySQL 5.6.22 - 5.7)	256M		
explicit-defaults-for-timestamp	0		

MySQL 5.5, 5.6 and 5.7			
Instance/Server Options			
sql-mode	See below ¹		
explicit_defaults_for_timestamp	0 (off/inactive) Deprecated in some releases		
innodb_file_per_table		1	
innodb_thread_concurrency		0	
table_open_cache		1000	
sort_buffer_size		2M	
read_buffer_size		2M	
tmp_table_size		400M	
max_heap_table_size		400M	
innodb_buffer_pool_size		4096M	
innodb_additional_mem_pool_size		20M	
binlog_format		row	
innodb_flush_log_at_trx_commit		2	
innodb_flush_method		O_DIRECT (Linux Only)	
innodb_doublewrite		0	
MySQL 5.5, 5.6 and 5.7			
Other Options			
Server Configuration Options	Defaults, unless instructed otherwise		

¹**sql-mode** must **not** be set to, or include one of the following:

- NO_ZERO_DATE
- NO_ZERO_IN_DATE
- "TRADITIONAL"

Examples of valid **sql-modes**:

- "ANSI"
- ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION

MySQL 5.5, 5.6 and 5.7			
Instance/Server Options			
	Mandatory	Recommended	Supported
[client]			
default-character-set	utf8		
[mysql]			
default-character-set	utf8		
[mysqldump]			
max_allowed_packet	250M		

- MySQL options may be written using underscore or hyphen in different contexts. Use the correct form according to your server version and usage context.

Manually Creating an OO Database on MySQL

During OO setup, a new database can be created automatically by the OO installer or a pre-existing database can be used.

If during installation, you are authorized to connect to the database server using the DBA role (connect as “root”), use the “create the database/schema” option, and you can skip this section.

This section describes the procedure for manually creating an OO database on MySQL.

Note: Only the database and user are created at this point; objects such as tables and indexes are not created yet. These objects are created once OO starts for the first time.

This section is relevant for you if, for example, due to security restrictions, you do not wish to use login credentials with elevated privileges during the OO installation. In such a case, you (or your organization’s DBA) should create the database first, and then let the OO installer connect to the pre-existing database using basic privileges.

To create a database, you must connect to the SQL Server instance using a user that has **CREATE** permission (at the very least).

- **root** has all privileges. Any member of the DBA role will also be able to create the user and database.

- Perform the following procedures only if you are an experienced MySQL database administrator.
- If you prefer to use the MySQL Workbench GUI, make sure you select all options that correspond with the SQL code presented below.
- Not all database creation options are specified—only those that differ from the default value. When in doubt, use default values.

To create a database:

1. Log in to the MySQL as “root” or any other member of the DBA role.
2. Edit and run `<installation folder>/central/bin/sql/MySQL/mysql_create_oo_db.sql` and verify that the database and user were created successfully. Adapt the script parameters to match your environment.

```
SET @OODB='OODB';
SET @OOUSER='OOUSER';
SET @OOPASS='OOPa55WorD';

SET @SQL1 = CONCAT('CREATE DATABASE IF NOT EXISTS `',@OODB,` COLLATE utf8_
bin');
SET @SQL2 = CONCAT('CREATE USER `',@OOUSER,`'@'%' IDENTIFIED BY
`',@OOPASS,`''');
SET @SQL3 = CONCAT('GRANT ALL PRIVILEGES ON `',@OODB,`.`.* TO
`',@OOUSER,`''');
PREPARE stmt1 FROM @SQL1;
PREPARE stmt2 FROM @SQL2;
PREPARE stmt3 FROM @SQL3;
EXECUTE stmt1;
EXECUTE stmt2;
EXECUTE stmt3;
DEALLOCATE PREPARE stmt1;
DEALLOCATE PREPARE stmt2;
DEALLOCATE PREPARE stmt3;
FLUSH PRIVILEGES;
```

3. Test your newly created connection to the database and verify that you are able to log in successfully.
4. (Optional) In order to verify that database objects can be created by the new login and user, connect to the database server using OO's newly created login role credentials and run:

`<installation folder>/central/bin/sql/MySQL/mysql_create_test_table.sql`

Verify that the script ran correctly and that no errors are shown.

```
USE OO;
```

```
CREATE TABLE TEST_TABLE(  
    TEST_COLUMN int NULL  
);  
  
INSERT INTO TEST_TABLE (TEST_COLUMN) VALUES ( 1 );  
INSERT INTO TEST_TABLE (TEST_COLUMN) VALUES ( 2 );
```

Manually Creating Database Objects

Once the database and user are in place, the database objects (tables, indexes, and so on) are created when the OO service starts and connects to the database for the first time.

You can create the database objects manually if the user created for OO does not have permissions to create or modify database objects (limited to DML operation only).

It is highly recommended to grant the OO database user DDL-related privileges, in order to allow OO to perform upgrade operations without manual intervention, as these operations sometimes require modifications to database structure.

To manually create the database objects:

1. Extract the **mysql.sql** file from the OO installation zip file under **docs\sql**.
2. Connect to the OO database as the relevant user, permitted to create and modify database objects in the OO database.
3. Run the **mysql.sql** file and verify that no errors occurred.

MySQL Database Maintenance

This section describes the various maintenance tasks that are recommended for OO databases created on MySQL, such as backing up the database, checking database integrity, handling index fragmentation, and monitoring the database.

This section includes:

Backing up the OO Database

You can back up MySQL database using several tools, such as **mysqldump** or **mysqlbackup**. You can back up the OO database using any type of method/tool as long as the complete database is backed up.

Consider the following guidelines when you create your backup plan for OO:

Backup method:

The backup method depends mainly on business considerations—how much information “may” be lost? What is the maximum time for system recovery? If you need to be able to perform point-in-time recovery, and only “allow” a few hours-worth of data loss, use the full recovery model and perform full or differential backups daily, and transaction log backup every N hours depending on your business requirements.

Backup frequency:

Daily backup is recommended, especially if you are using/modifying OO on a daily basis.

You should back up once a month at the very least.

Timing:

Schedule backup for the time when OO is least active.

Retention:

Retention depends on your business guidelines and regulations.

Creating a Maintenance Plan

Maintaining an OO database includes rebuilding the index and reclaiming free space. Use the scripts and tools described in this section, in order to keep the OO database in good shape.

Recommended Utility for Database Maintenance

In order to keep the OO database in good shape, it is recommended to schedule **mysqlcheck** utility to run during a system maintenance window.

Important! Note that this operation locks tables! Only perform this action during a maintenance window when the OO system is not operating!

Here is an example of how to run this utility:

```
mysqlcheck -uouser -p???? -os --auto-repair OO
```

Replace "oouser" and "OO" with the actual OO user name and database name, respectively.

It is recommended not to provide the password explicitly. See the MySQL documentation for recommendations on how to secure database passwords.

Utility for Purging Historical Data

To install and use the OO history purging stored procedure:

- Download the latest **HPE_OO_DB_Maintenance.zip** pack from ITOM Marketplace and unpack it to a temporary folder. Follow the instructions in *OO Database Maintenance Procedures* to install and execute a purge stored procedure for MySQL.

Upgrade/Rollback Instructions

Prior to any upgrade or rollback of OO, take the following steps:

1. Purge non-essential data from the database.

Upgrade and rollback operations convert flow execution data. Less data means a shorter upgrade/rollback. If all flow execution data is essential, perform a complete database backup first and then purge.
2. Back up your database, preferably after the database was purged.
3. Schedule the operation to the database server's quiet period.
4. Verify that no shared resources are exhausted – temporary space, transaction log, and similar resources.
5. Make a backup copy of your MySQL connector jar file. Consider using a newer one for upgrade scenarios.

Following any upgrade or rollback of OO, take the following steps:

1. Install relevant versions of maintenance and purge stored procedures.

Stored procedures and purge flows are not upgraded/rolled back by the OOupgrade utility.
2. Reschedule database maintenance and purge routines (stored procedures and/or purge flows).

Deploying and Maintaining a Postgres Database

In order to deploy (OO) using Postgres, you must have an existing Postgres database service. If you need to create a new database service, see the relevant documentation provided by Postgres, because this information is not included within this section. However, this section contains recommendations for the Postgres configuration.

This chapter includes the following sections:

Workflow for Postgres Deployment

To deploy OO using Postgres, perform the following steps:

1. **Review sizing guidelines.** For details, see "OO Database Sizing" in ["Introduction to Preparing the Database Environment" on page 28.](#)
2. **Review Hardware and Software Requirements.** For details, see ["System Requirements for Postgres" below.](#)
3. **Configure Postgres.** For details, see ["Configuring Postgres" on page 69.](#)
4. **Create OO database on Postgres.** For details, see ["Manually Creating an OO Database on Postgres" on page 70.](#)

System Requirements for Postgres

This section describes the system requirements for working with Postgres in conjunction with OO.

Hardware Requirements

For OO database sizing guidelines and hardware requirements, see "OO Database Sizing" and "Hardware Requirements" in ["Introduction to Preparing the Database Environment" on page 28.](#)

For Postgres hardware requirements, see the relevant installation section for your Postgres release and operating system.

Software Requirements

The following table lists the Postgres releases supported by OO:

Postgres Database Releases	
Version	Type
9.1	x86 32-bit
	x86 64-bit
9.2	x86 32-bit
	x86 64-bit
9.3	x86 32-bit
	x86 64-bit
9.4	x86 32-bit
	x86 64-bit
9.5	x86 32-bit
	x86 64-bit
9.6	x86 32-bit
	x86 32-bit

Only supported versions should be used.

See the Postgres documentation for supported platforms.

Examples of Tested Deployments

The following table lists the deployment environments that have been rigorously tested by the quality assurance personnel.

Database Release		Operating System
Version	32/64Bit	
Postgres 9.2.3	64 Bit	Windows 2012 Standard Edition 64 Bit
Postgres 9.1.9	64 Bit	Red Hat Enterprise Linux 6.3 64 Bit
Postgres 9.3.2	64 Bit	Red Hat Enterprise Linux 6.3 64 Bit
Postgres 9.4.6	64 Bit	Red Hat Enterprise Linux 7 64 Bit
Postgres 9.5.1	64 Bit	Red Hat Enterprise Linux 7 64 Bit

Language Support

Postgres determines character set and collation at the database level. The OO database uses Unicode (utf8) encoding and collation. This lets you use any Unicode character (and practically all common characters in all languages).

Configuring Postgres

This section contains information on Postgres configuration settings.

Legend:

- **Mandatory** configuration options/values appear in **bold/orange** font.
- **Recommended** configuration options/values appear in **bold/purple** font.
- Supported configuration options/values appear in normal font, and may show as a comma-separated list.
- *Comments* appear in *italic* font.

Postgres 9.1 - 9.6			
Instance/Server Options			
Instance Configuration Options	Defaults, unless instructed otherwise		
	Mandatory	Recommended	Supported
max_connections	>= 1000		
default_transaction_isolation	'read committed'		
autovacuum	on		
track_counts	on		
shared_buffers	>=512MB ¹		
effective_cache_size	>=2048MB ¹		
work_mem	>=1MB ¹		
maintenance_work_mem	>=32MB ¹		
lc_messages		'en_US.UTF-8'	Any
lc_monetary		'en_US.UTF-8'	Any

[1] - Minimal values. See the Postgres documentation on how to tune these values in accordance with your environment.

Manually Creating an OO Database on Postgres

During OO setup, a new database can be created automatically by the OO installer or a pre-existing database can be used.

If during installation, you are authorized to connect to the database server as a privileged user (connect as “postgres”), use the “create the database/schema” option, and you can skip this section.

This section describes the procedure for manually creating an OO database on Postgres.

Note: Only the database and role are created at this point; objects such as tables and indexes are not created yet. These objects are created once OO starts for the first time.

This section is relevant for you if, for example, due to security restrictions, you do not wish to use login/user credentials with elevated privileges during the OO installation. In such a case, you (or your organization’s DBA) should create the database, login, and user first, and then let the OO installer connect to the pre-existing database using basic privileges.

To create a database, you must connect to the Postgres instance using a login that has **CREATEUSER** and **CREATEDB** privileges at the very least.

- The **postgres** built-in user has all the required privileges.
- Perform the following procedures only if you are an experienced Postgres database administrator.
- If you prefer to use the PgAdmin GUI, make sure you select all options that correspond with the SQL code presented below.
- Not all database creation options are specified—only those that differ from the default value. When in doubt, use default values.

To create a database:

1. Log in to Postgres as “postgres” or any other login role with **CREATEUSER** and **CREATEDB** privileges.
2. Edit and run:

```
<installation folder>/central/bin/sql/PostgreSQL/postgres_create_oo_db_linux.sql
```

or:

postgres_create_oo_db_windows.sql

Verify that the database, login, and user were created successfully. Adapt the script parameters to match your environment.

```
CREATE ROLE "oouser" LOGIN
UNENCRYPTED PASSWORD '???????'
NOSUPERUSER INHERIT NOCREATEDB NOCREATEROLE NOREPLICATION;

CREATE DATABASE "OO"
WITH OWNER = "oouser"
ENCODING = 'UTF8'
TABLESPACE =
LC_COLLATE =
LC_CTYPE =
CONNECTION LIMIT = 1000;
```

3. (Optional) In order to verify that database objects can be created by the new login and user, connect to the database server using OO's newly created login role credentials and run:

<installation folder>/central/bin/sql/PostgreSQL/postgres_create_test_table.sql

Verify that the script ran correctly and that no errors are shown.

```
CREATE TABLE TEST_TABLE(
    TEST_COLUMN int NULL
);

INSERT INTO TEST_TABLE (TEST_COLUMN) VALUES ( 1 );
INSERT INTO TEST_TABLE (TEST_COLUMN) VALUES ( 2 );
```

Manually Creating Database Objects

Once the database and role are in place, the database objects (tables, indexes, and so on) are created when the OO service starts and connects to the database for the first time.

To manually create the database objects (instead of the OO service):

1. Extract the **postgres.sql** file from the OO installation zip file under **docs\sql**.
2. Connect to the OO database as the OO database user.
3. Run the **postgres.sql** file and verify that no errors occurred.

Postgres Database Maintenance

This section describes the various maintenance tasks that are recommended for OO databases created on Postgres, such as backing up the database, checking database integrity, handling index fragmentation, and monitoring the database.

This section includes:

Backing up the OO Database

You can back up a Postgres database using several tools, such as the **pg_dump** or **pg_backup** script. You can back up the OO database using any type of method/tool as long as the complete database is backed up.

Consider the following guidelines when you create your backup plan for OO:

Backup method:

The backup method depends mainly on business considerations—how much information “may” be lost? What is the maximum time for system recovery? If you need to be able to perform point-in-time recovery, and only “allow” a few hours-worth of data loss, use the full recovery model and perform full or differential backups daily, and transaction log backup every N hours depending on your business requirements.

If your organization is more tolerant to data loss, you can use the simple recovery model and perform a daily or weekly full backup.

Backup frequency:

Daily backup is recommended, especially if you are using/modifying OO on a daily basis.

You should back up once a month at the very least.

Timing:

Schedule backup for the time when OO is least active.

Retention:

Retention depends on your business guidelines and regulations.

Creating a Maintenance Plan

OO Postgres database maintenance mainly includes table REINDEX, as **autovacuum** needs to be activated. Use the example below, in order to keep the OO database in good shape.

Recommended Utility for Database Maintenance

In order to keep the OO database in good shape, it is recommended to run the REINDEX action during a system maintenance window.

Important! Note that this operation locks tables! Only perform this action during a maintenance window when the OO system is not operating!

Here is an example of how to REINDEX a complete database using the **reindexdb** utility:

```
reindexdb -d OO -U oouser -w ?????
```

Replace “OO” and “oouser” with actual OO database and user names.

It is recommended not to provide the password explicitly. See the Postgres documentation for recommendations on how to secure database passwords.

Utility for Purging Historical Data

OO 10.x keeps all flow and step execution data in the database by default, enabling comprehensive debugging of previous flow runs. As a result, the database size grows in accordance to the system throughput and flow complexity. It is highly recommended to track your database size and make sure that old, irrelevant information is periodically purged.

Reclaiming free space in Postgres database requires two phases:

1. Space is marked as deleted – following a DELETE command.
2. A background “vacuum” process runs to allow free space to be reused.

Note that after the execution of the purge procedure described below, the vacuum process will usually spring into action and start “vacuuming” tables and indices.

It is important to allow the vacuum process to complete successfully before starting another purge. If the purge and vacuum overlap, they stall each other as they compete for object locks.

“Vacuum Full” Operation

PostgreSQL database is only able to reclaim Large Objects (LOBs) space when the **Vacuum Full** operation is run. A considerable amount of the OO database space is stored as LOBs. Therefore, it is important to schedule a routine **Vacuum Full** operation - at times in which OO is least active. The **Vacuum Full** operation requires *exclusive table locks*, and therefore may interfere with activity of OO.

To install and use the OO history purging stored procedure:

1. Download the latest **HPE_OO_DB_Maintenance.zip** pack from ITOM Marketplace and unpack it to a temporary folder. Follow the instructions in the *OOAdministration Guide* ["Set Up Database](#)

[Environment" on page 28](#) to install and execute a purge stored procedure for PostgreSQL.

```
SELECT OOPurgeHistory(90,10000,0,1,4,0);
```

Upgrade/Rollback Instructions

Prior to any upgrade or rollback of OO, take the following steps:

1. Purge non-essential data from the database.

Upgrade and rollback operations convert flow execution data. Less data means a shorter upgrade/rollback. If all flow execution data is essential, perform a complete database backup first and then purge.

2. Verify that you have sufficient disk space to support a "vacuum full" operation, and initiate a "vacuum full" operation.
3. Back up your database, preferably after the database was purged.
4. Schedule the operation to the database server's quiet period.
5. Verify that sufficient free disk space is available for transaction log files.

Following any upgrade or rollback of OO, take the following steps:

1. Install the relevant version of the purging stored procedure.

Stored procedures and purge flows are not upgraded/rolled back by the OO upgrade utility.

2. Reschedule the database purge routines (stored procedure and/or purge flows).
3. Reschedule the routine database "vacuum full" operations.

Database Purging

Purge Flows

Purge flows are provided as content in the HPE Solutions content pack, and are the preferred method for purging the OO database. As long as purge flows provide sufficient "throughput" (are able to delete a sufficient amount of rows in a given amount of time), it is recommended to use them instead of the purge stored procedures.

Purge flows have several advantages over stored procedures:

- Manageability – The OO administrator can easily schedule and track purge operations.
- Database privileges – No need for assistance from the DBA in installing and scheduling purge operations.
- Purge functionality upgrades – When OO is upgraded, and a new table requires purging, this capability will be added to an existing flow or a new one will be created for that purpose.

New stored procedures, however, must be manually downloaded and installed.

Note: For each purge type (rerun, steplog, and so on), you can only run one purge flow at a time.

Purge Best Practices

Relating to purge operations, it is important to distinguish between two separate criteria:

- Retention – the age of the data you wish to keep
- Purge frequency – how often the purge operations are executed

These two criteria are completely unconnected. You can keep one years worth of data, while purging every 30 minutes.

As the retention period is often determined as a tradeoff between database size and business requirements, the purge frequency can vary between hourly, daily, or even weekly. The recommended purge frequency is between “hourly” and “daily”.

Purging smaller amounts of data frequently proves (in most cases) to be a better strategy than purging large amounts of data during a weekly or monthly maintenance window.

Here is a suggested workflow for optimizing purge frequency:

- Schedule a recurring purge operation for every four hours and monitor it:
 - How long does it take for the purge operation to complete (average and max. duration)?
 - What was the effect on database performance?
 - What was the effect on Operations Orchestration’s performance (user interface and flow execution)?
- As long as purge operations complete in under 30 minutes and no dramatic effect was registered by the DBA or the database, keep this frequency.

You may also try frequencies from hourly to daily and see which works best.

Note that some effect on the database and OO is considered normal.

- Otherwise, if OO's activity changes considerably throughout the day, and consistent "quiet periods" exist, schedule purge operations to run during these periods.

Additional Guidelines for Microsoft SQL Server

This appendix contains additional guidelines relevant for a DBA's deployment on Microsoft SQL Server.

Using Windows Authentication to Access Microsoft SQL Server Databases

Unless configured otherwise, DBAs use Microsoft SQL Server authentication to access Microsoft SQL Server databases. Note that the OO installer currently does not support using Windows authentication during the OO installation. However, Windows authentication can be used once OO is installed.

This appendix describes how to enable OO to use Windows authentication to access Microsoft SQL Server databases.

Configuring OO to Work with Windows Authentication

You can enable OO to use Windows authentication instead of Microsoft SQL Server authentication to access OO databases.

To enable OO to use Windows authentication to access a Microsoft SQL database:

1. Encrypt the Windows user password using the **encrypt-password.bat** utility located under **<OO installation>/central/bin** by running:

```
encrypt-password.bat --encrypt --password <password>
```

Save the generated string in order to use it in the next step.

2. Back up your current **database.properties** file located under **<OO installation>/central/conf** if you have an existing (usable) database connection.

3. Edit the **database.properties** file located under **<OO installation>/central/conf**, and change only the relevant parameter syntax to match the following example.

Note: When copying scripts, note that you may need to remove redundant line breaks from the copied version.

For SQL Server 2008R2 and 2012:

```
db.username=[USERNAME]  
db.password=[the string generated by encrypt-password.bat]
```

```
jdbc.url=jdbc\:jtds\:sqlserver\://[DB_HOSTNAME]\:[PORT]/[DB_NAME]  
;sendStringParametersAsUnicode=true  
;domain\[DOMAIN_NAME]
```

For SQL Server 2014:

```
db.username=[USERNAME]  
db.password=[the string generated by encrypt-password.bat]
```

```
jdbc.url=jdbc\:sqlserver\://[DB_HOSTNAME]\:[PORT];databaseName\[DB_NAME];sendStringParametersAsUnicode=true;integratedSecurity=true;
```

For SQL server 2014, there might be a need to copy **sqljdbc_auth.dll** to the system path of the central server machine. For more details, see the Microsoft JDBC documentation.

Replace all the highlighted items with the correct values that match your environment.

Note that the **jdbc.url** parameter is broken in this document into several lines for readability. When copying this example onto your **database.properties** file, remove all line breaks so the **jdbc.url** parameter is composed of a single line containing no white spaces.

Configuring OO to Work with Always On

For SQL Server 2008R2 and 2012

OO connects to SQL Server 2008R2 and 2012 databases using the jtds connector. Thus, it is only able to register a single IP address for the Availability Group Listener. In addition, OO does not support Multi-Subnet Failover.

Cluster Configuration

Configure your cluster environment as follows:

1. Set the cluster option **RegisterAllProvidersIP** to 0 for the network name associated with the Availability Group Listener.
2. Optionally, decrease the **HostRecordTTL** value for the same cluster resource to reduce database reconnection time.

See the Microsoft SQL Server Documentation (“Create or Configure an Availability Group Listener”) for a complete description of configuration options and examples.

Installing Operations Orchestration using Availability Group Listener

Prior to installing OO, verify that the Availability Group Listener network name resolves to an accessible IP address from all Central machines.

1. In the OO installation wizard **Database Connection Configuration** page, enter the Availability Group Listener network name in the **Hostname or IP address** field.
2. Enter the rest of the details and click the **Test Connection** button.

Providing the Availability Group Listener network name guarantees that the OO Central server(s) will be able to reconnect to the database in the case of a failure.

For silent installation, use:

```
db.url=jdbc:jtds:sqlserver://[AG-NET-NAME]:[PORT]/[DB-NAME];sendStringParametersAsUnicode=true
```

For SQL Server 2014

OO connects to SQL Server 2014 databases using the Microsoft JDBC 4.2 connector.

Here is an example on how to format JDBC URL in order to connect to an Always On cluster:

```
db.url=jdbc:sqlserver://[AG-NET-NAME];instanceName=[NAMED-INST-NAME];  
databaseName=[DB-NAME];multiSubnetFailover=true;applicationIntent=ReadWrite;  
sendStringParametersAsUnicode=true
```

Note that **multiSubnetFailover** and **applicationIntent** are added in order to allow more complex behaviors.

Note that Multi-Subnet Failover scenarios in conjunction with SQL Server 2014 Always On cluster were not certified with OO.

For additional information and connection options, refer to "JDBC Driver Support for High Availability, Disaster Recovery" in the Microsoft documentation and "Setting the Connection Properties" in the Microsoft JDBC documentation.

Additional Guidelines for Oracle

This appendix contains the configuration that needs to be done for OO to work with Oracle 11gR2 and 12cR1 Real Application Cluster (RAC) environments. This information is for advanced users only.

This appendix includes:

Oracle Real Application Cluster (RAC)

A cluster is a collection of interconnected servers that appear as one server to the end user and to applications. Oracle Real Application Cluster (RAC) is Oracle's solution for high availability, scalability, and fault tolerance. It is based on clustered servers that share the same storage.

Oracle RAC is a single Oracle database installed on a cluster of hardware servers. Each server runs an instance of the database and all the instances share the same database files.

For more details about Oracle RAC, see the Oracle Clusterware Guide and the Oracle Real Application Clusters Administration and Deployment Guide in the Oracle documentation set of your release.

In this appendix, the following Oracle RAC example is used:

- Oracle RAC Cluster name: OORAC
- Service Name: ORCL.MY.DOMAIN
- Machine names: Server1, Server2
- On each machine, there is an Oracle instance of OORAC:
 - SID on Server1: OORAC1
 - SID on Server2: OORAC2
- On each machine, there is a virtual IP (Server1-Vip and Server2-Vip):
 - Server1-Vip is assigned to Server1
 - Server2-Vip is assigned to Server2

The virtual IP is in addition to the static IP assigned to the machine.

- SCAN listener uses a virtual IP, usually published using DNS/GNS:

SCAN-Vip

- The local listeners on both servers listen on the default port 1521 and support the database service OORAC. SCAN listener(s) reside on one of the cluster nodes, and may failover along with their

virtual IP addresses in the case of a failure.

Note: Although 12c SCAN allows configuring multiple SCAN listeners (per subnet), OO only supports connecting to a single 12c SCAN listener.

Single Client Access Name

In release 11g, Oracle introduced the Single Client Access Name (SCAN), as a preferred access method for clients connecting to the RAC. In this method, clients are not required to configure individual nodes in the RAC; rather, they use a single virtual IP known as the SCAN or the SCAN VIP.

The SCAN is a single network name defined for the cluster either in your organization's Domain Name Server (DNS) or in the Grid Naming Service (GNS) that rotates between several IP addresses, reflecting multiple listeners in the cluster. The SCAN eliminates the need to change clients when nodes are added to or removed from the cluster.

The SCAN and its associated IP addresses provide a stable name for clients to use for connections, independent of the nodes that make up the cluster. 11g SCAN addresses, virtual IP addresses, and public IP addresses must all be on the same subnet. 12c SCAN allows configuring multiple SCAN listeners – one per subnet. OO only supports connecting to a single 12c SCAN listener.

The SCAN method is recommended when using OO in an Oracle 11g RAC environment.

Configuring OO to Work with Oracle RAC

Connecting to a SCAN Listener Virtual IP

1. In the OO installation wizard **Database Connection Configuration** page, enter the SCAN listener virtual IP address/network name in the **Hostname or IP address** field.
2. Select the **Service Name** radio button instead of the **SID** radio button and enter the Oracle RAC service name.
3. Enter the rest of the details and click the **Test Connection** button.

Providing the SCAN listener virtual IP address guarantees that the OO Central server(s) will be able to reconnect to the database cluster in the case of a failure.

This connection method is the preferred one as it enables the OO installation wizard to create the database as well as populate it.

For silent installation use (replace the highlighted text with the actual values):


```
db.url=jdbc:oracle:thin:@//[SCAN-Vip]:[PORT]/[ORCL.MY.DOMAIN]
```

Load-Balancing Using Explicit Connection String

Although using Oracle's SCAN listener virtual IP is the preferred method, you can also provide an explicit connection string by using the **Other database** installation option. For complete details, see "Oracle RAC Example" in "[Installation Wizard 'Other database' Option](#)" below.

Here is an example of how a load balancing connection string would appear in Central's **database.properties** file.

1. Back up your current database.properties file located under **<OO installation>/central/conf** if you have an existing (usable) database connection.
2. Edit the **database.properties** file as follows, while replacing the highlighted items with the values that match your environment.

Note: When copying scripts, note that you may need to remove redundant line breaks from the copied version.

```
jdbc.url=jdbc\:oracle\:thin\:@  
(DESCRIPTION=(LOAD_BALANCE=on)(ADDRESS_LIST=  
(ADDRESS=(PROTOCOL=TCP)(HOST=Server1-Vip)(PORT=1521))  
(ADDRESS=(PROTOCOL=TCP)(HOST=Server2-Vip)(PORT=1521))))  
(CONNECT_DATA=(SERVICE_NAME=ORCL.MY.DOMAIN)))
```

Note that the **jdbc.url** above is broken into several lines for readability, while it should appear as a single line in the actual configuration file.

When Load Balancing is set to "on", failover between listeners is active by default.

Installation Wizard 'Other database' Option

This appendix contains information about the **Other database** option in the OO installation wizard.

This option enables using specific JDBC driver and connection options. Use this option if:

- You want to use a different version of the JDBC driver, other than the ones provided in the OO installation (see the note below).
- You want to provide a JDBC connection URL yourself, to include an option or options which are not currently provided by the standard database connection options.

NOTES:

- Connection is restricted to the types and versions of database described in the "Overview" section in ["Introduction to Preparing the Database Environment" on page 28.](#)
- Use a non-provided JDBC driver at your own risk. The following JDBC drivers are supported:
 - jtds-1.3.0.jar – for SQL Server 2008R2 and 2012
 - sqjjdbc4-4.2.jar – for SQL Server 2014
 - ojdbc7-12.1.0.2.jar
 - postgresql-9.4.1207.jar
 - mysql-connector-java-5.1.35 (not provided in the OO installation)

When you install OO using the **Other database** option, the installation wizard does not create the database and/or related user or roles. You must create these beforehand, using the installation wizard. Detailed instructions on how to create the database/user/role can be found in this document, under the "Manually Creating an OO Database" section in each database chapter.

Microsoft SQL Server Named Instance Example

SQL Server 2008R2 and 2012

The following example show a connection to a Microsoft SQL Server Named Instance using jtds JDBC Connector. The database, login role, and user are pre-created by the DBA, and the login role is the owner of the database (has full DML and DDL privileges).

Use the following details in the installation wizard, and revise the highlighted values to match your environment.

Use JDBC URL Option #1 when a unique TCP port is assigned to each instance.

Use JDBC URL Option #2 when **Database Browser :service** is active and is able to route connections based on the instance name. Note that TCP port number does not appear in this option.

Field	Value	Comments
JDBC Driver jar	C:\my\path\jtds-1.3.0.jar	
JDBC Driver class name	net.sourceforge.jtds.jdbc.Driver	
JDBC URL	jdbc:jtds:sqlserver://<DB_IP_OR_HOSTNAME>:<DB_	

Field	Value	Comments
Option #1	PORT>/<DB_NAME>;instance=<INSTANCE_NAME>;sendStringParametersAsUnicode=true	
JDBC URL Option #2	jdbc:sqlserver://<DB_IP_OR_HOSTNAME>:<INSTANCE_PORT>;databaseName=<DB_NAME>;sendStringParametersAsUnicode=true	
Username	<LOGIN_ROLE>	
Password	<LOGIN_ROLE_PASSWORD>	

SQL Server 2014

The following example show a connection to a Microsoft SQL Server Named Instance using Microsoft JDBC Connector. The database, login role, and user are pre-created by the DBA, and the login role is the owner of the database (has full DML and DDL privileges).

Use the following details in the installation wizard, and revise the highlighted values to match your environment.

Use JDBC URL Option #1 when a unique TCP port is assigned to each instance.

Use JDBC URL Option #2 when Database Browser service is active and is able to route connections based on the instance name. Note that TCP port number does not appear in this option.

Field	Value	Comments
JDBC Driver jar	C:\my\path\sqljdbc4-4.2.jar	
JDBC Driver class name	net.sourceforge.jtds.jdbc.Driver	
JDBC URL Option #1	jdbc:sqlserver://<DB_IP_OR_HOSTNAME>:<INSTANCE_PORT>;databaseName=<DB_NAME>;sendStringParametersAsUnicode=true	INSTANCE_PORT represents the separate instance
JDBC URL Option #2	jdbc:sqlserver://<DB_IP_OR_HOSTNAME>;instanceName=<INSTANCE_NAME>;databaseName=<DB_NAME>;sendStringParametersAsUnicode=true	If Database Browser service is active
Username	<LOGIN_ROLE>	
Password	<LOGIN_ROLE_PASSWORD>	

Microsoft SQL Server Windows Authentication Example

SQL Server 2008R2 and 2012

The following example shows a connection to a Microsoft SQL Server using Windows Authentication and jtds JDBC connector. The database is pre-created by the DBA, and the Windows login account is the owner of the database (has full DML and DDL privileges).

Use the following details in the installation wizard, and revise the highlighted values to match your environment:

Field	Value	Comments
JDBC Driver jar	C:\my\path\jtds-1.3.0.jar	Use jtds JDBC driver only
JDBC Driver class name	net.sourceforge.jtds.jdbc.Driver	
JDBC URL	jdbc:jtds:sqlserver://<DB_IP_OR_HOSTNAME>:<DB_PORT>/<DB_NAME>;domain=<DOMAIN_NAME>;sendStringParametersAsUnicode=true	
Username	<WINDOWS_USERNAME>	
Password	< WINDOWS_USERNAME_PASSWORD >	

Note: Ignore the error regarding incompatible collation, but make sure to set your collation to one of the supported collations described in the “Language Support” section in ["Introduction to Preparing the Database Environment" on page 28](#).

It is recommended but not required, that OO Central service(s) is run as the same user that is used for database authentication. Note that using a domain account for authentication may enforce security policies that may change from time to time by the domain administrator.

Whenever the domain account password changes, the new password must be encrypted and the new encrypted password must be placed in the **database.properties** file located under **<OO installation>/central/conf**. If an OO cluster is used, this operation must be repeated for all OO Central servers.

SQL Server 2014

The following example show connection to a Microsoft SQL Server 2014 using Windows Authentication and Microsoft JDBC Connector.

Prerequisites:

1. The database is pre-created by the DBA
2. The relevant Windows account is the owner of the database (has full DML and DDL privileges)
3. OO's installation wizard must be unzipped and extracted to a folder form
4. **sqljdbc_auth.dll** must be downloaded and copied to OO's installation wizard under **<installer_folder>\java\bin**. This dll file is not provided along with the OO installation and is available for download from Microsoft web site. Download Microsoft JDBC 4.2 and extract 64bit version of **sqljdbc_auth.dll**.

Installation steps:

1. Run the OO installer by executing **<installer_folder>\installer.bat** as the same Windows account that owns the OOdatabse.
2. In the **Connectivity** page, select the **Do not start central server after installation...** checkbox.
3. In the **Database Connection** page select the **Other database** option and fill in the details according to the table below. Edit and replace highlighted parts with actual values.
4. After the installation:
 - a. Verify **sqljdbc_auth.dll** is located in **<central installation>\java\bin**.
 - b. Configure the service to run under the same Windows account that owns the OO database.
 - c. Start the OO service and verify successful connection to the database.

Repeat the steps above for all central server installations

Use the following details in the installation wizard, and revise the highlighted values to match your environment:

Field	Value	Comments
JDBC Driver jar	C:\my\path\sqljdbc4-4.2.jar	
JDBC Driver class name	com.microsoft.sqlserver.jdbc.SQLServerDriver	
JDBC URL	jdbc:sqlserver://<DB_IP_OR_HOSTNAME>:<DB_PORT>;databaseName=<DB_NAME>;integratedSecurity=true;sendStringParametersAsUnicode=true	
Username	Fully qualified Windows account: DOMAIN\USERNAME	
Password	---- keep empty ----	

Note: Ignore the error about incompatible collation, but make sure to set your collation to one of the supported collations described in the “Language Support” section in ["Introduction to Preparing the Database Environment" on page 28.](#)

Note that using a domain account for authentication may enforce security policies that may change from time to time by the domain administrator.

Oracle RAC Example

The following example shows a connection to an Oracle RAC cluster using explicit JDBC URL. The database (user) must be pre-created by the DBA, and must be granted DML and DDL privileges.

Use the following details in the installation wizard, and revise the highlighted values to match your environment:

Field	Value	Comments
JDBC Driver jar	C:\my\path\ojdbc7-12.1.0.2.jar	
JDBC Driver class name	oracle.jdbc.OracleDriver	
JDBC URL	jdbc:oracle:thin:@(DESCRIPTION=(LOAD_BALANCE=on) (ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=Server1_VIP)(PORT=1521)))(ADDRESS= (PROTOCOL=TCP)(HOST=Server2_VIP)(PORT=1521))) (CONNECT_DATA=(SERVICE_NAME=RAC1.MY.DOMAIN)))	
Username	<PRE-CREATED-DATABASE>	
Password	< PRE-CREATED-DATABASE-PASSWORD>	

Securing and Hardening OO

This document explains how to deploy and manage OO instances in a secure manner, and how to configure security hardening for OO.

Introduction

This guide is designed to help IT professionals who deploy and manage (OO) instances in a secure manner. Our objective is to help you make well-informed decisions about the various capabilities and features that OO provides to meet modern enterprise security needs.

Security requirements for the enterprise are constantly evolving and this section should be viewed as the best effort to meet those stringent requirements. If there are additional security requirements that are not covered by this section, please open a support case with the support team to document them and we will include them in future editions of this section.

This section is designed to help IT professionals who deploy and manage Operations Orchestration (OO) instances in a secure manner. Our objective is to help you make well-informed decisions about the various capabilities and features that OO provides to meet modern enterprise security needs.

Security requirements for the enterprise are constantly evolving and this section should be viewed as 's best effort to meet those stringent requirements. If there are additional security requirements that are not covered by this section , please open a support case with the support team to document them and we will include them in future editions of this section.

Technical System Landscape

OO is an enterprise-wide application based on Java 2 Enterprise Edition (J2EE) technology. J2EE technology provides a component-based approach to the design, development, assembly, and deployment of enterprise applications.

Security Updates

Between OO 10.20 and 10.50 the following security updates were made:

Between OO version 10.20 and 10.50 the following security updates were made:

- When the **Enable Capture of Logged-in User Credentials** check box is selected in Central, OO will temporarily capture (in a secure manner) the credentials of the logged-in user when this user runs flows in the Remote Debugger. A message warns that credentials may be captured.
- In OO 10.5x, the default is that there is no default role. This gives the administrator better control of user authorization, because users only get roles that are explicitly assigned either to them or to their LDAP group.
- When OO has multiple LDAP configurations, if the administrator flags one of these as default, users who belong to it will not be required to select a domain upon login.

- OO 10.5x secures sensitive data (for example, passwords) during the execution. If a variable is marked as sensitive in Studio, it will be retrieved in an encrypted form when being used in scriptlets.

Between OO 10.10 and 10.20 the following security updates were made:

- It is now possible to grant permissions for system accounts in OO. This enables the administrator to control which users can view which system accounts and run flows that use them. This feature is useful for customers with multiple organizations, who may wish to hide some of the system accounts from some users.
- It is now possible to apply permissions to multiple roles in the Edit Permissions dialog box. In previous versions, it was only possible to select one role at a time.
- When you upgrade an OO installation from an earlier 10.x version, the SSL truststore is updated to include the up-to-date trusted root certificates, as published by Oracle. This includes deletion of expired certificates, and import of new ones.
- OOnow offers the option to audit events, so that you can track security breaches. Auditing lets you track actions that took place on Central, such as logins, triggering flows, creating schedules, editing configurations, and so on.

Currently, you can retrieve the audit trail only via APIs.

- OOnow supports encryption keys that are 2048 bits long (and longer). This aligns our cryptography keys with the FIPS 186-4 standard.
- A new `sslEnabledProtocols` property has been added to the `server.xml` file (located at `<installation_folder>/central/tomcat/conf/server.xml`):

```
sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2"
```

This property ensures that only TLS v1, TLS v1.1 and TLS v1.2 are allowed and that SSL 3.0 is not. This prevents vulnerability to the “POODLE” attack (Padding Oracle On Downgraded Legacy Encryption).

Security Overview

This section provides an overview of the security models and recommendations for a secure implementation of OO. This includes subjects such as authentication, authorization, encryption, and more. Where relevant, there are references to other OO documents, which describe how to complete security-related tasks.

Security Concepts

OO Glossary

For more information about OO concepts, see the *OO Concepts Guide*.

Role Permission

A permission is a predefined authorization to perform a task. OO Central includes a set of permissions that can be assigned to [roles](#).

For example, the **Schedule** permission grants the ability to view and create run schedules.

Role

A role is a collection of [permissions](#).

For example, the **Flow Administrator** role may be assigned the **View Schedules** permission and the **Manage Schedules** permission.

User

A user is an object associated with a person (or application identity) representing the person and defining their authorization.

[Roles](#) are assigned to users, to define the actions they are authorized to perform in Central. For example, the user Joe Smith may be assigned the **Flow Administrator** role.

It is possible to configure different kinds of users:

- **LDAP users** log on to Central using their LDAP user name and password. For example, using their Active Directory user name and password.
- **Internal users** log on to Central using the user name and password that was set up locally in Central.
- **LWSSO** - Lightweight Single Sign On (SSO) is a mechanism in which a single action of user authentication and authorization can permit a user to access all systems that support LWSSO. For

example, if users have logged onto another product web client that has LWSSO enabled, they can enter the OO Central application directly, bypassing the OO Central logon screen.

When an internal user and an LDAP user with the same role are logged in, there is no difference between their permissions.

Note: It is recommended to use LDAP users rather than internal users, because LDAP users are secured according to policies implemented by the LDAP provider.

Content Permission

Content permission is permission to view or run individual flows or the flows in a particular folder.

Users who have been assigned a specified role will be able to access the flows according to the content permissions assigned to their role.

For example, users with the **Administrator** role may be entitled to view and run all the flows in the system, while users with the **User** role may be entitled to run certain flows, and have view permission for others.

Common Security Concepts

System Security

The processes and mechanisms by which computer-based equipment, information, and services are protected from unintended or unauthorized access, change, or damage.

Least Privilege

The practice of limiting access to the minimal level that will allow normal functioning. This means giving a user account only those privileges that are essential to that user's work.

Authentication

The process of identifying an individual, usually based on a user name and password, or certificate.

Authorization

Permission to access system objects, based on an individual's identity.

Encryption

A way to enhance the security of a message or file by scrambling the contents so that it can be read only by someone who has the right encryption key to encode it. For example, the TLS protocol encrypts the communication data.

Countermeasure

A way to mitigate the risk of a threat.

Defense in Depth

Layers of protection, so that you don't have to rely on a single security measure alone.

Risk

A possible event that could cause damage. For example, financial loss, damage to the company image, and so on.

Threat

Triggering a risk event that exploits a vulnerability.

Vulnerability

A weakness in a target that can potentially be exploited by a security threat.

Secure Implementation and Deployment

OO Security Hardening

The "Hardening" chapter provides recommendations for safeguarding your OO deployment from security risks or threats. Some of the most important reasons to secure an application include protecting the confidentiality, integrity, and availability of an organization's critical information.

To comprehensively protect your OO system, it is necessary to secure both OO and the computing environment (for example, the infrastructure and the operating system) upon which the application runs.

The "Hardening" chapter provides recommendations to help secure OO at the application level and does not cover how to secure the infrastructure within the customer environment. The customer is solely responsible for understanding his/her infrastructure/environment and applying the respective hardening policies.

Physical Security

It is recommended that OO is protected by physical security controls defined by your organization. The OO server components are installed in a physically secured environment, according to best practice. For example, the server must be in a closed room with access control.

Secure Installation Guidelines

Supported Operating Systems

For the types and versions of supported operating systems, see the *OO System Requirements*.

Operating System Hardening Recommendations

Contact your operating system vendor for recommended best practices for hardening your operating system.

For example:

- Patches should be installed
- Unnecessary services/software should be removed or disabled
- Minimal permissions should be assigned to users
- Auditing should be enabled

Tomcat Hardening

When you install OO Central, Tomcat is partially hardened by default. If you want extra hardening, see the recommendations in the Hardening chapter.

Installation Permissions

The following permissions are required to install and run OO:

Installing OO	Windows/Linux: Any standard user who is able to run a Java process, and who has permission to create folders and services
Running OO	<ul style="list-style-type: none">• Windows: The Windows service runs as the system user or a specific user (the user must have access to the OO installation directory)• Linux: Any standard user who is able to run a Java process

See also the recommendations in the CIS Apache Tomcat documentation.

Network and Communication Security

The *OO Architecture Guide* describes the basic OO topology, high availability, and load balancer security.

The *OO Network Architecture White Paper* describes the required firewall configuration, and suggests two workarounds that can be applied in cases when, due to policy restrictions, the required firewall configuration cannot be implemented:

- SSH reverse tunneling
- Reverse proxy

Communication Channel Security

Supported Protocols and Configuration

OO supports the TLS protocol.

For more information, see ["Replacing the Central TLS Server Certificate" on page 110](#).

The Central ports are defined by the administrator during the installation.

Channel Security

OO supports the following secure channels:

Channel (Directed)	Supported Secure Protocol
OOSH, browser, Studio Remote Debugger, or RAS → Central	For a secure channel, use the TLS communication for encryption and Client Certificate for authentication.
Central → LDAP Server	To encrypt communication between Central and LDAP, use Secure LDAP, using the TLS protocol.

RAS Security

In a topology with a reverse RAS (which waits for Central to initiate the connection), the following mechanism protects the security of the RAS:

- If there are several consecutive failed connection attempts (because the wrong shared secret was entered), this will result in a delay.

For more information about reverse RASes, see the "Setting Up Topology – Workers and RASes" in the *OO Central User Guide*.

Administration Interface Security

Accessing the Administration Interface

There are several ways to control access to the administration interface:

- Credentials
- Client certificate
- SAML

Securing the Administration Interface - Recommendations

1. Authentication must be enabled in Central.
See "Enabling Authentication" in the *OO Central User Guide*
2. It is recommended to secure the administration interface with the TLS protocol. You should set up TLS between the client and the Central interface for encryption.
See "[Working with Server and Client Certificates](#)" on page 109.
3. It is recommended to work with LDAP users, rather than internal users, because this is more secure.
4. It is recommended to set up authentication to access Central via Client Certificates. This is more secure than user passwords.
See "[Working with Server and Client Certificates](#)" on page 109.

User Management and Authentication

Authentication Model

To enable easy bootstrapping of the authentication mechanism in OO, the product starts with authentication disabled.

You must enable authentication immediately after installation.

For information about how to enable authentication, see "Enabling Authentication" in the *OO Central User Guide*.

There are a number of ways to authenticate access to Central.

Choose the method of identifying users:

- User name and password
- Client Certificate
- SAML token
- Single sign on (LWSSO)

Choose one of two ways to manage the users:

- LDAP users , saved on an LDAP server as Active Directory (recommended)
- Internal users and passwords, saved locally on the Central server (not recommended)

Types of Users

Different types of users can have different permissions assigned for them. For example, flow author, administrator, system administrator, and so on.

Authentication Administration and Configuration

Internal or LDAP Users

You can set up internal users with passwords in the Central UI or define the user in the LDAP server and map LDAP groups to Central roles.

Note: Our recommendation is not to use internal users, but to use a more secure alternative such as LDAP users.

Database Authentication

OO supports four databases: Oracle, MS SQL, MySQL, and Postgres.

We recommend using a strong database password for database authentication and using a strong password policy. For example, blocking after a number of failed attempts.

When using MS SQL, it is possible to work with either database authentication or with OS authentication. Our recommendation is to work with OS authentication, where this is possible. For example, it is possible to use Windows authentication to access Microsoft SQL Server databases.

Authorization

Authorization Model

User access to OO resources is authorized based on the user's role, and the permissions configured for that role.

See:

- "Setting Up Security – Roles" in the *OO Central User Guide*
- "Assigning Permissions to a System Account" in the *OO Central User Guide*.

Minimal Permissions Guidelines

It is recommended to:

- Select appropriate permissions for the role.
- Use minimal permissions when creating new roles.
- Grant minimal permissions and extend the permissions only as needed to avoid unwanted privilege escalation. For example, start with View permissions and add additional permissions individually as needed.

Authorization Configuration

Central is installed with a number of out-of-the box roles, which you can configure and assign to users. By default, the out-of-the box roles are assigned the following permissions:

Role	Default Permissions
Administrator	All
End_user	None
Everybody	None
Promoter	All the Content permissions
System_admin	All the System permissions

Default Role

It is possible to configure one of the roles with the **Default Role** attribute. If you do so, make sure that this is the role with the least privileges. Remember that when you give permissions to this role, this affects all LDAP users, in addition to those are explicitly associated with the role.

For more information, see "Assign a role to be the default role" under "Setting Up Security – Roles" in the *OO Central User Guide*.

See also:

- "Assigning Permissions to a System Account" in the *OO Central User Guide*.
- "Setting Permission for Content" in the *OO Central User Guide*.

Access to Workspaces in Studio

When creating multiple workspaces in Studio, we recommend creating a workspace under folders to which only the creating user has read and write permissions.

Workspaces created under public folders might be accessible to all users, making them prone to tampering and disclosure of sensitive information.

Configure OO Central in IDM Mode

When Central is connected to IDM, authentication is enabled by default at startup. To use OO Central in the IDM mode, you must switch from the LWSSO security profile to the IDM security profile.

Switch to the IDM security profile

1. Open the `<installation dir>\central\conf\central-wrapper.conf` file in a text editor and locate the `wrapper.java.additional` parameter called `-Doverride.startup.mode` under the `# Java Additional Parameters` section.

2. Update the value of the `-Doverride.startup.mode` parameter to `IDM`. For example:

```
wrapper.java.additional.33=Doverride.startup.mode=IDM
```

To switch back to the LWSSO security profile, update the value of the `-Doverride.startup.mode` parameter to `LWSSO`.

3. Save and close the `central-wrapper.conf` file.
4. Restart the Central service.

Update the `idm.properties` file

You can configure IDM properties in the `<installation dir>\central\conf\idm.properties` file. After configuring a property, save and close the file. Then restart the Central service.

The following table describes the properties in the `idm.properties` file.

Property	Description
<code>idm.configuration.url</code>	Represents the URL of the IDM service. For example, <code><idm_protocol>://<idm_hostname>:<idm_port>/<idm_service_path></code>

Property	Description
idm.configuration.username	User name of the REST IDM integration
idm.configuration.password	Password of the REST IDM integration. This value can be encrypted using the <installation dir>\central\bin\encrypt-password.bat tool.
idm.configuration.internal.username	User name of the internal account (used for PAS authentication and API calls towards the IDM service). This user has the IDM_ADMIN permission to be able to perform IDM administration operations for consumer organizations.
idm.configuration.internal.password	Password of the internal account (used for PAS authentication and API calls towards the IDM service). This value can be encrypted using the <installation dir>\central\bin\encrypt-password.bat tool.
idm.configuration.scheduler.username	User name of the scheduling account used to execute Central flow schedules.
idm.configuration.scheduler.password	Password of the scheduling account used to execute Central flow schedules. This value can be encrypted using the <installation dir>\central\bin\encrypt-password.bat tool.
idm.configuration.signing.key	Signing key of the IDM configuration
idm.configuration.oo.central.tenant	Tenant value of Central configured in the IDM service. The default value is 00_Central.
idm.configuration.oo.central.return.url	Return URL of Central after being redirected from the IDM login page. For example, <code><central_protocol>://<central_hostname>:<central_port>/oo</code> For more information, see "Configure the return URL of Central in IDM mode" below.

Configure the return URL of Central in IDM mode

When Central is running in IDM mode, users trying to visit Central are first redirected to the IDM login page to provide their credentials. This return URL is used to redirect the user back to Central's dashboard after a successful login. The value of the return URL can be configured using the `idm.configuration.oo.central.return.url` property in the `idm.properties` file.

Note: Ensure that the protocol and ports used in the return URL are consistent throughout the entire login flow. For example, you cannot start the login flow through HTTP and provide an HTTPS return URL.

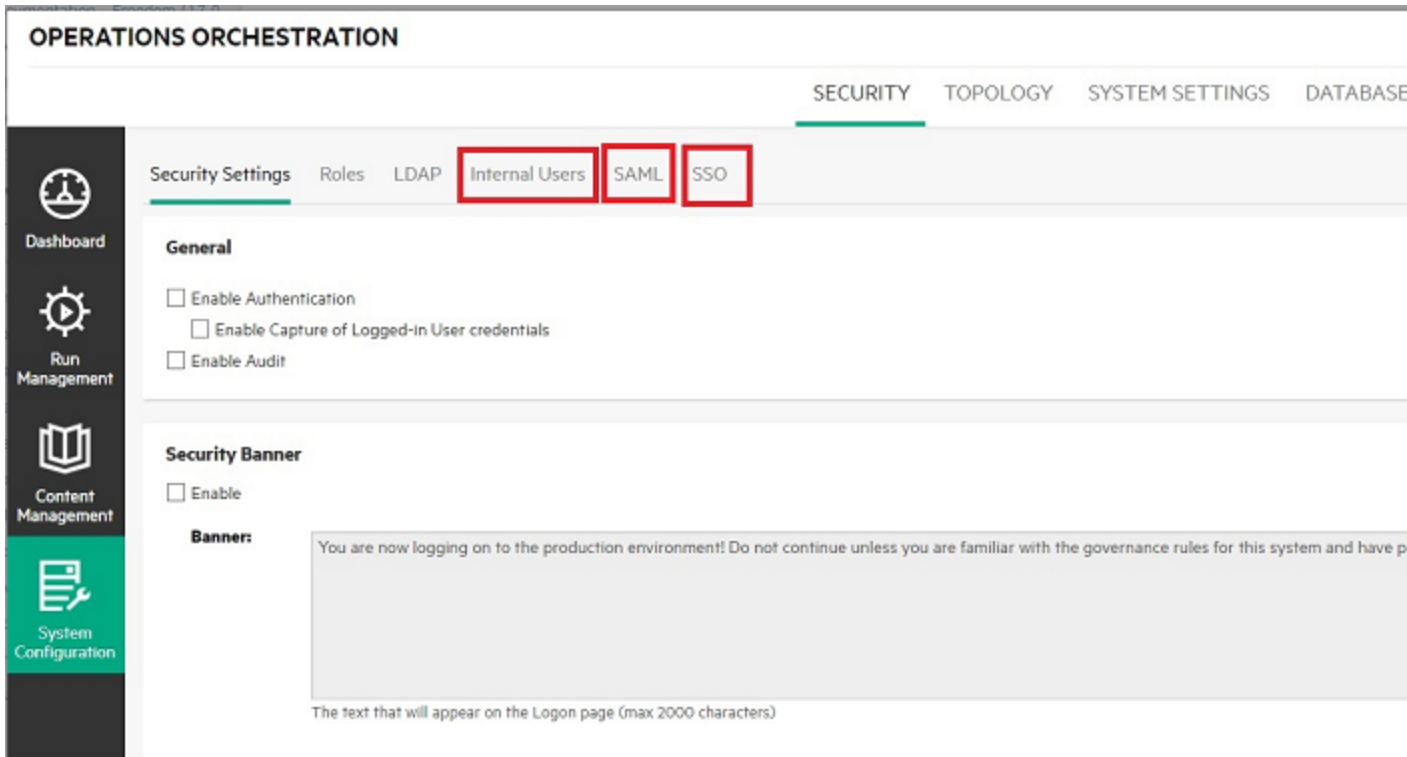
Example of the `idm.properties` file

```
# IDM Properties
idm.configuration.url=https://idm-hostname:8443/idm-service/
idm.configuration.username=idmTransportUser
idm.configuration.password={ENCRYPTED}uNEBrRr29t/78P700j3BiA==
idm.configuration.internal.username=admin
idm.configuration.internal.password={ENCRYPTED}30l15qHg/7770l15qHgP700d2MiA==
idm.configuration.signing.key={ENCRYPTED}uWRUrRQ4BzUNr2EG/4PDfB700j9LiA==
idm.configuration.oo.central.tenant=OO_Central
idm.configuration.oo.central.return.url=https://central-hostname:8433/oo
```

Hidden Tabs in the IDM Mode

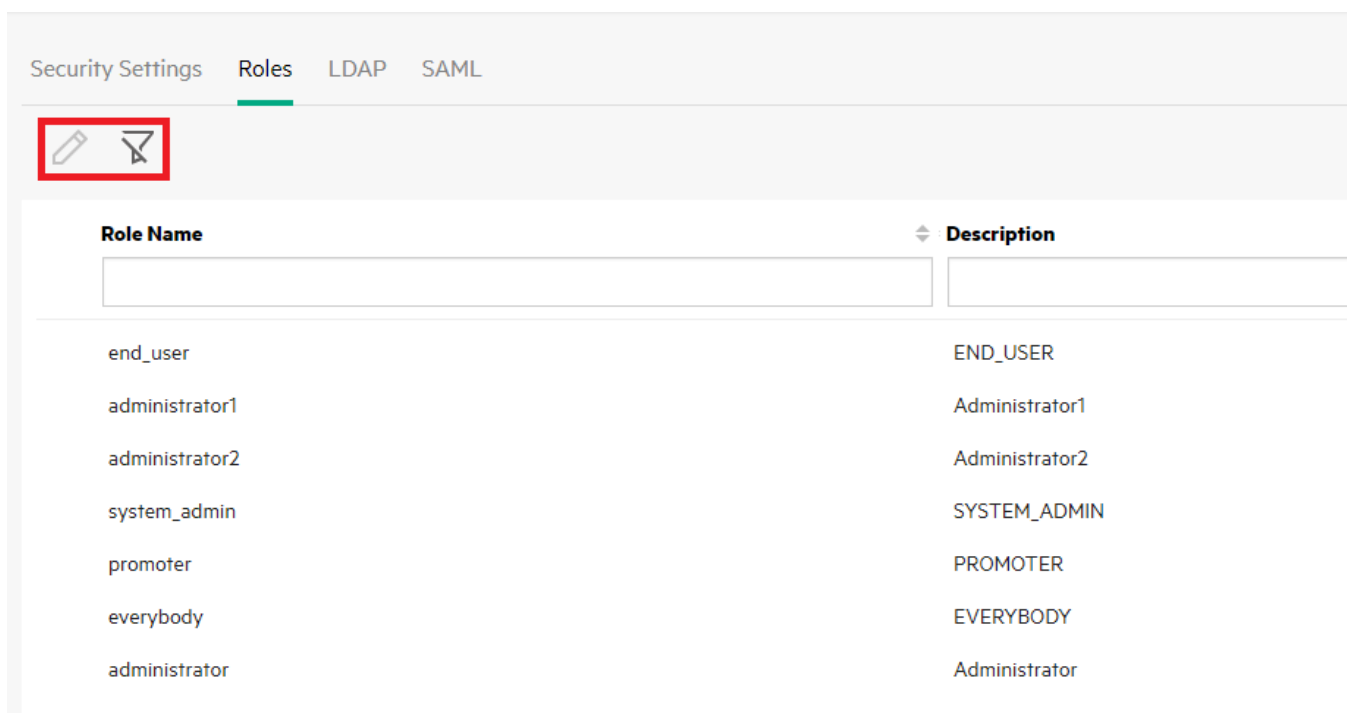
OO authentication allows the OO administrator to define internal users; however, in the IDM mode, the OO administrator cannot define internal users. Therefore, some tabs are hidden when OO is used in the IDM mode.

The following figure shows the tabs that are hidden in the IDM mode.





Roles in IDM

In the IDM mode, the OO administrator cannot create additional roles. The administrator can map only existing roles to groups. Therefore, the **Roles** tab in the UI displays only the **Edit** and **Clear Filters** icons.



Security Settings Roles LDAP SAML

Role Name	Description
end_user	END_USER
administrator1	Administrator1
administrator2	Administrator2
system_admin	SYSTEM_ADMIN
promoter	PROMOTER
everybody	EVERYBODY
administrator	Administrator

Configure LDAP in IDM Mode

When Central is configured to use IDM for authentication, you can add an LDAP configuration for IDM directly from Central.

The IDM server will use the LDAP configuration to authenticate and authorize user access to Central. Central supports a single LDAP configuration for an IDM server.

Add an LDAP configuration for the IDM server

1. Click the **System Configuration** button to display the **System Configuration** workspace.
2. Select **Security > LDAP**.
3. Click **Add**.
The LDAP tab displays the fields relevant to the LDAP configuration.
4. In the **Base DN** field, enter the root DN for fetching users.
5. In the **User ID** field, enter the distinguished name of a user with search privileges.
It is recommended to provide a privileged user, because some HPE OO functionality uses searching, for example, the scheduler. If there is no privileged user, this functionality may not work

properly, depending on how the LDAP is configured.

For example, `uid=john,ou=people,dc=devlab,dc=ad`

Ensure that you use the exact DN for the user with these privileges.

6. In the **Password** field, enter the password of the user with search privileges.
7. In the **Connections** section, select **Secured Channel** to support secured channel transport.
8. Click **Add** in the **Connections** section to add a the connectivity details of the LDAP host.
9. In the **Host** field, enter the IP address or host name of the LDAP server.
10. In the **Port** field, enter the port number of the LDAP server. This value must be between 0 and 64435.
11. In the **User Name Attribute** field, enter the LDAP attribute that contains the user name.
12. In the **Common Name** field, enter the attribute that defines how the user name will be displayed on the screen. In most LDAP implementations, **cn** is the default.
13. In the **User Search Base** field, enter the root DN for fetching users.
14. In the **Users Search Filter** field, enter the search filter parameters in order to filter users. This filter will be applied on the user DN that was entered in the **User Search Base** field.
15. In the **User Email** field, enter the LDAP attribute that contains the user email.
16. Select the **Users SUB Search Scope** check box if you want the users filter to run recursively.
17. In the **Group Membership** field, enter the LDAP attribute that defines to which groups the user belongs.
18. In the **Manager Identifier** field, enter the LDAP attribute that identifies the group.
19. In the **Manager Identifier Value** field, enter the search value for the **Manager Identifier** field.
20. Click **Save** to save the configuration.

If the LDAP configuration is not saved, an asterisk appears next to the domain name in the LDAP navigation pane to the left. If you leave the page before saving, a message appears, warning that there are unsaved changes.

Backup

In order to prevent data loss, It is highly recommended to regularly back up your data on the servers onto secure media. This is also helpful for disaster recovery and business continuity.

After installing OO, make sure to back up the **central\var\security** folder and the **central\conf\database.properties** file.

Some data on the database schema is encrypted and the keys for decryption are stored locally on the OO Central server. If these system files become corrupted or deleted, the schema will be useless, because there will be no way to decrypt the data.

Note: The keys are encrypted, so it is important to include them in the backup. The keys are located in the **security** folder.

See:

- "Backing Up OO" in the *OO Administration Guide*
- "Setting up Disaster Recovery" in the *OO Administration Guide*
- "Backing Up and Recovering the Central Security Files" in the *OO Installation, Upgrade, and Configuration Guide*
- "Using a Load Balancer in OO Deployment" in the *OO Architecture Guide*

Encryption

Encryption Model

OO supports encryption and hash algorithms to protect sensitive data. Encryption is designed to prevent the exposure and modification of sensitive data, such as passwords, definitions, and so on, in OO.

It is important to use well-known, standard algorithms without known vulnerabilities, in order to prevent decryption by unauthorized persons.

For example, SSL is not used, because of known vulnerabilities in the SSL protocol.

Static Data

All saved passwords are protected using well known algorithms and none appear in clear text.

For example:

- The system account passwords are encrypted.
- The internal user passwords are hashed.
- The database passwords are encrypted.

Data in-transit

OO uses the Transport Layer Security (TLS) protocol to encrypt the data between components (such as Central and RAS).

Disabling the HTTP port

It is recommended to disable the HTTP port, for security reasons, so that the only communication channel will be on TLS and encrypted. For more information, ["Changing the HTTP/HTTPS Ports or Disabling the HTTP Port" on page 121](#).

Encryption Administration

Recommended Encryption Best Practices

In order to reach higher levels of security and cryptography, it is recommended to configure OO to be compliant with Federal Information Processing Standards (FIPS) 140-2. OO can be set to be compliant with FIPS 140-2 Level 1.

Default Configuration Set

- Symmetric-key algorithm: AES with key size 128
- Hashing algorithm: SHA1

Advanced Settings

After you have configured OO for FIPS 140-2 compliance, OO uses the following security algorithm:

- Symmetric-key algorithm: AES256
- Hashing algorithm: SHA256

See ["Configuring Compliance with FIPS 140-2" on page 136](#).

Digital Certificates

A digital certificate is an electronic "passport" for a person, server, station, and so on.

- To use encryption between a browser and the Central server, you need to install a digital certificate on the server side.
- To use Client Certificate to authenticate the Central server, you need to install a Client Certificate on the client side (for example, on the browser, RAS, OOSH, Studio, and so on).

OO uses the Java Keytool utility to manage cryptographic keys and trusted certificates. This utility is included in the OO installation folder, in **<installation dir>/java/bin/keytool**.

Certificate Location

Installations of OO Central include two files for the management of certificates using Keytool:

- **<installation dir>/central/var/security/client.truststore**: Contains the list of trusted certificates
- **<installation dir>/central/var/security/key.store**: Contains the OO private certificate (including the private key)

Access Control to KeyStore and TrustStore

It is recommended that the TrustStore and KeyStore are stored with read permissions only for the user that runs the Central service.

Replacing the OO Self-signed Certificate

It is recommended to replace the OO self-signed certificate after a new installation of OO or if your current certificate has expired.

Part of the process of replacing the certificate is generating a PKCS12 format certificate, using your CA. Contact your CA for specific details about the certificate process or refer to your corporate policy.

For more information, see ["Replacing the Central TLS Server Certificate" on page 110](#).

Adding Digital Signatures to a Content Pack

If a content pack has a digital signature from a trusted CA, this provides assurance that the content is trusted.

Adding a digital signature is not mandatory.

- OO out-of-the-box content packs contain a digital signature from Verisign.
- OO authors are recommended to add a digital signature to their custom content packs.
- If a signed content pack is breached, the content pack cannot be deployed.
- If the signature is expired, a warning appears before deployment, and the user must select a check box acknowledging that they are ignoring the expired signature.

Pay attention to content packs that are not signed. An unsigned content pack is not trusted, and could contain malicious content. Note also that unsigned content pack could be breached and with the signature removed.

For more information about digital certification of content packs, see "Deploying and Managing Content Packs" in the *OO Central User Guide*.

Sensitive Information in a Content Pack

System Account Passwords

Do not include passwords when creating a content pack. The passwords will be obfuscated inside the content pack, which is not a secure option.

The OO security best practice is to configure the system account passwords in Central. For more information, see "Setting Up System Accounts for a Content Pack" in the *OO Central User Guide*.

Hardening OO

This section describes how to configure security hardening for OO.

Note: For information about other administrative tasks, see the *OO Installation, Upgrade, and Configuration Guide*.

Security Hardening Recommendations

1. Install the latest version of OO. For more information, see the *OO Installation, Upgrade, and Configuration Guide*.
2. (Optional) Configure OO for FIPS 140-2 Compliance. If you choose to do this, it must be configured before you start the Central server. See "[Configuring OO for FIPS 140-2 Level 1 Compliance](#)" on page 134.
3. Configure the Central server certificate for TLS encryption and client certificate for strong authentication (mutual).

Note: This can be done during installation.

For the RAS, Debugger, and OOSH, provide certificate authentication if required (for the server certificate) and use the client certificate for authentication against the Central. See "[Working with Server and Client Certificates](#)" on the next page.

4. Harden the OO Central server by removing the HTTP port and replacing the passwords of the KeyStore and TrustStore with strong passwords. See "[Changing the HTTP/HTTPS Ports or Disabling the HTTP Port](#)" on page 121 and "[Changing and Encrypting/Obfuscating the](#)

[KeyStore/TrustStore Password" on page 117.](#)

5. Harden OO Studio by replacing the KeyStore and TrustStore passwords with strong passwords, and encrypt or obfuscate the passwords in the configuration files. See ["Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password" on page 117.](#)
6. Remove the RC4 cipher from the SSL-supported ciphers. See ["Removing Vulnerable Ciphers from the SSL supported Ciphers" on page 121.](#)
7. (Optional) Configure the TLS protocol version. See ["Configuring the TLS Protocol" on page 140.](#)
8. Enable authentication in Central. See "Enabling Authentication" in the *OO Central User Guide*.
Internal users are not secured, so use a secured LDAP with a strong password policy. See "Setting Up Security – LDAP Authentication" in the *OO Central User Guide*.
9. Harden/secure the operating system and database.
10. Add a security banner with a meaningful message. For example, "You are now logging on to our PRODUCTION environment! Do not continue unless you are familiar with the governance rules for this system and have taken the necessary training." See "Setting Up a Security Banner" in the *OO Central User Guide*.
11. In the Windows and SQL server environment, configure OO to work with Windows authentication. See "Configuring OO to Work with Windows Authentication" in the *OO Database Guide*.
12. Make sure that auditing is enabled in Central. For more information, see "Enabling Auditing" in the *OO Central User Guide*.

Working with Server and Client Certificates

Transport Layer Security (TLS) certificates digitally bind a cryptographic key to the details of an organization, enabling secure and encrypted connections from a web server to a browser.

OO uses the Keytool utility to manage cryptographic keys and trusted certificates. This utility is included in the OO installation folder, in `<installation dir>/java/bin/keytool`. For more information about the Keytool utility, see

<http://docs.oracle.com/javase/7/docs/technotes/tools/solaris/keytool.html>.

Note: Keytool is an open source utility.

Installations of OO Central include two files for the management of certificates:

- `<installation dir>/central/var/security/client.truststore`: Contains the list of trusted certificates.
- `<installation dir>/central/var/security/key.store`: Contains the OO certificate (private key).

Note: If you use client certificates with LDAP, the LDAP must be configured as default in Central. For details, see "Setting Up Security – LDAP Authentication" in the *OO Central User Guide*

Recommendations:

- It is recommended to replace the OO self-signed certificate after a new installation of OO or if your current certificate is expired.
- It is recommended to store the TrustStore and KeyStore with read permissions only for the user that runs the Central service.
- It is recommended to clear the console after using Keytool or to use the prompt for password inputs.

Encrypting the Communication Using a Server Certificate

Replacing the Central TLS Server Certificate

You can use a certificate signed by a well-known certificate authority or a custom server certificate from a local certificate authority.

Replace the parameters that are highlighted in **<yellow>** to match the location of the **key.store** file and other details on your computer.

Note: The following procedure uses the Keytool utility that is located in **<installation dir>/java/bin/keytool**.

1. Stop Central and back up the original **key.store** file, located in **<installation dir>/central/var/security**.
2. Open a command line in **<installation dir>/central/var/security**.
3. Delete the existing server certificate from the Central **key.store** file, using the following command:

```
keytool -delete -alias tomcat -keystore key.store -storepass <keystore password>
```

4. If you already have a certificate with a **.pfx** or **.p12** extension, go to the next step. If not, you need to export the certificate with private key into PKCS12 format (.pfx,.p12). For example, if the certificate format is PEM:

```
>openssl pkcs12 -export -in <cert.pem> -inkey <.key> -out <certificate name>.p12 -name <name>
```

If the certificate format is DER, add the `-inform DER` parameter after `pkcs12`. For example:

```
>openssl pkcs12 -inform DER -export -in <cert.pem> -inkey <.key> -out  
<certificate name>.p12 -name <name>
```

Note:

To generate the PKCS12 format certificate, you need to use your CA. As this step may vary according to CA vendor and policy, contact your CA for a detailed explanation of the certificate generation process.

Note: Make a note of the password that you provide. You will need this password for the private key when you input the KeyStore passphrase later in this procedure.

Make sure to choose a strong password.

5. List the alias for your certificate's alias, using the following command:

```
keytool -list -keystore <certificate_name> -v -storetype PKCS12
```

The certificate alias is displayed and should be provided in the next command.

In the example below, it is the fourth line from the bottom.

```
C:\Program Files\Hewlett-Packard\oo-saml\central\var\security>keytool -list -keystore server.pfx -v -storetype PKCS12  
Enter keystore password:  
Keystore type: PKCS12  
Keystore provider: SunJSSE  
Your keystore contains 1 entry  
Alias name: 1e-775fb32c-269c-499b-bae8-fe7077479ec6  
Creation date: 24/04/2014  
Entry type: PrivateKeyEntry  
Certificate chain length: 2
```

6. Import the PKCS12 format server certificate to the Central **key.store** file using the following command:

```
keytool -importkeystore -srckeystore <PKCS12 format certificate path> -  
destkeystore key.store -srcstoretype pkcs12 -deststoretype JKS -alias <cert  
alias> -destalias tomcat
```

7. If the imported server certificate has a different password from the original server certificate, it is important to change the `keyPass` password. Follow the instructions in ["Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password"](#) on page 117.

It is also recommended to change the default "changeit" password in the automatically generated KeyStore in the Central server. See ["Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password"](#) on page 117.

8. Start Central.

Importing a CA Root Certificate to the Central TrustStore

If you are using a custom root certificate for Central, you will need to import the trusted root certificate authority (CA) to the **client.truststore**. If you are using a well known root CA (like Verisign) you do not have to perform the following procedure, because the certificate will already be in the **client.truststore** file.

By default, OO supports all self-signed certificates. However, in a production environment, it is recommended to change this default to a custom CA or a well-known CA for security reasons.

Replace the parameters that are highlighted in **<yellow>**.

Note: The following procedure uses the Keytool utility that is located in **<installation dir>/java/bin/keytool**.

1. Stop Central and back up the original **client.truststore** file, located in **<installation dir>/central/var/security/client.truststore** or **<oosh dir>/var/security/client.truststore** for a standalone OOSH.
2. Import the trusted root certificate authority (CA) to the Central **client.truststore** file if it doesn't already exist in the CA list (by default, all the well known CAs are there). In a standalone OOSH, it is under the main OOSH directory.

```
keytool -importcert -alias <any_alias> -keystore <path to the  
client.truststore> -file <certificate_name.cer> -storepass <changeit>
```

3. Start Central.

Importing a CA Root Certificate to a RAS TrustStore

After installing a RAS, if you are using a custom root certificate for Central and you didn't provide this root certificate during the RAS installation, you will need to import the trusted root certificate authority (CA) to the RAS **client.truststore**. If you are using a well known root CA (like Verisign) you do not have to perform the following procedure, because the certificate will already be in the **client.truststore** file.

By default, OO supports all self-signed certificates. However, in a production environment, it is recommended to change this default to a custom CA or a well known CA for security reasons.

Replace the parameters that are highlighted in **<yellow>**.

Note: The following procedure uses the Keytool utility that is located in **<installation dir>/java/bin/keytool**.

1. Stop the RAS and back up the original **client.truststore** file, located in **<installation dir>/ras/var/security/client.truststore**.
2. Open command line in **<installation dir>/ras/var/security**.
3. Open the **<installation dir> ras/conf/ras-wrapper.conf** file and make sure that the `Dssl.support-self-signed` value is set to **false**. This enables the trusted root certificate authority (CA).

For example:

```
wrapper.java.additional.<x>=-Dssl.support-self-signed=false
```

4. Open the **<installation dir> ras/conf/ras-wrapper.conf** file and make sure that the `Dssl.verifyHostName` is set to **true**. This verifies that the FQDN in the certificate matches the FQDN of the request.

For example:

```
wrapper.java.additional.<x>=-Dssl.verifyHostName=true
```

Note: This property is set to **true** by default.

5. Import the trusted root certificate authority (CA) to the RAS **client.truststore** file if it doesn't already exist in the CA list (by default, all the well known CAs are there):

```
keytool -importcert -alias <any_alias> -keystore <path to the client.truststore> -file <certificate_name.cer> -storepass <changeit>
```

6. Start the RAS.

Importing a CA Root Certificate to the OOSH TrustStore

If you are using a custom root certificate for Central, you will need to import the trusted root certificate authority (CA) to the OOSH **client.truststore**. If you are using a well known root CA (like Verisign) you do not have to perform the following procedure, because the certificate will already be in the **client.truststore** file.

By default, OO supports all self-signed certificates. However, in a production environment, it is recommended to change this default to a custom CA or a well known CA for security reasons.

Replace the parameters that are highlighted in **<yellow>**.

Note: The following procedure uses the Keytool utility that is located in **<installation dir>/java/bin/keytool**.

1. Stop Central and back up the original **client.truststore** file, located in **<installation dir>/central/var/security/client.truststore** or **<oosh dir>/var/security/client.truststore** for a standalone OOSH.
2. Edit the **oosh.bat** from **<installation dir>/central/bin** or the main OOSH directory for a standalone OOSH.
3. Make sure that the **-Dssl.support-self-signed** value is set to **false**. This enables the trusted root certificate authority (CA).

For example:

```
-Dssl.support-self-signed=false
```

4. Make sure that the **-Dssl.verifyHostName** is set to **true**. This verifies that the FQDN in the certificate matches the FQDN of the request.

For example:

```
-Dssl.verifyHostName=true
```

Note: This property is set to **true** by default.

5. Import the trusted root certificate authority (CA) to the Central **client.truststore** file if it doesn't already exist in the CA list (by default, all the well known CAs are there):

```
keytool -importcert -alias <any_alias> -keystore <path to the  
client.truststore> -file <certificate_name.cer> -storepass <changeit>
```

6. Run OOSH.
7. Start Central.

Importing a CA Root Certificate to the Studio TrustStore

If you are using custom certificates in the Central, SVN, or GIT servers, in order for Studio to be able to work with these, you will need to import the trusted root certificate authority (CA) to the Studio **client.truststore** file. If you are using a well known root CA (like Verisign) you do not have to perform the following procedure, because the certificate will already be in the **client.truststore** file.

By default, OO supports all self-signed certificates. However, in a production environment, it is recommended to change this default to a custom CA or a well known CA for security reasons.

For a fresh **.oo** folder, Studio copies the **client.truststore** file from **<installation.dir>/studio/var/security** to the **<user>/.oo** folder. This is a one time action, in order to ensure that Studio can automatically import certificates (for example, for the Studio Remote

Debugger). Studio will use this file as the **client.truststore** if it exists; otherwise, it will use the one from the Studio installation (`<installation.dir>/studio/var/security/client.truststore`).

After an upgrade to 10.5x or later, the truststore location is the `<user>/.oo` folder.

If you want to manually import a certificate, you can import it either to `.oo/client.truststore` or to **client.truststore** in the Studio installation folder.

If you are using multiple workspaces, the changes made to the **client.truststore** file located under the `.oo` folder will apply only to the specific workspace. In order to apply the change to all the newly created workspaces, edit the **client.truststore** file located in the Studio installation folder.

Note: The following procedure uses the Keytool utility that is located in `<installation dir>/java/bin/keytool`.

1. Close Studio and back up the original **client.truststore** file, located in `<user>/.oo`
For example, `C:/Users/<username>/.oo`
2. Edit the **Studio.I4j.ini** file from `<installation dir>/studio`.
3. Make sure that the `-Dssl.support-self-signed` value is set to **false**. This enables the trusted root certificate authority (CA).

For example:

```
-Dssl.support-self-signed=false
```

4. Make sure that `-Dssl.verifyHostName` is set to **true**. This verifies that the FQDN in the certificate matches the FQDN of the request.

For example:

```
-Dssl.verifyHostName=true
```

5. Import the trusted root certificate authority (CA) to the Studio **client.truststore** file if it doesn't already exist in the CA list (by default, all the well known CAs are there). Replace the parameters that are highlighted in **<yellow>**:

```
keytool -importcert -alias <any_alias> -keystore <path to the  
client.truststore> -file <certificate_name.cer> -storepass <changeit>
```

6. Start Studio.

For more information, see the User Guide .

Checking the Revoke Status of a Certificate

A certificate revocation list (CRL) is a list of certificates (or more specifically, a list of serial numbers for certificates) that have been revoked. Entities that present these (revoked) certificates should no longer be trusted.

From the RAS Side

The RAS can determine when a Central certificate has been revoked in the remote server. When the RAS is started and performs a handshake with Central, it retrieves the Central certificate that contains a link to the location of the CRL file. The RAS then goes to the CRL file and validates the Central certificate against this CRL file.

If a certificate has been revoked, there is no connection between the RAS and the Central. In addition, there will be error messages in the log file on the RAS side.

From the Central Side

On the Central side, the RAS appears with Availability Offline (not connected) in the Topology area.



What do you want to do?

Enable Revoke Status Check

1. Open the RAS wrapper file **ras-wrapper.conf** located under **<RAS_INSTALLATION>/raslconf**.

2. Add the following lines to the RAS:

```
wrapper.java.additional.<n>=-Dcom.sun.security.enableCRLDP=true  
wrapper.java.additional.<n>=-Dcom.sun.net.ssl.checkRevocation=true
```

3. Set the following flag to false:

```
ssl.support-self-signed=false
```

Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password

Changing the KeyStore, TrustStore, and Server Certificate Passwords in the Central Configuration

1. Make sure that Central is running.

Note: Before doing this step, make sure that there are encrypted passwords. For information about how to encrypt a password, see "Encrypting Passwords" in the *OO Administration Guide*.

From OOSH, run the following command:

```
set-sys-config --key <keyName> --value <encryptedPassword>
```

where <keyName> is one of the values from the table below:

Configuration item	Action
<code>key.store.password</code>	To set the password used to access to the key.store . The default value is "changeit". This needs to correspond with the value for <code>keystorePass</code> , as set in the steps below.
<code>key.store.private.key.alias.password</code>	To set the password used to access the server certificate (private key) from the key.store . The default value is "changeit". This needs to correspond with the value for <code>keyPass</code> , as set in the steps below.

2. Stop the Central service.
3. Change the KeyStore, TrustStore, and server certificate password using Keytool.

Use the following keytool command to change the KeyStore password:

```
keytool -storepasswd -keystore <installation_  
folder>/central/var/security/key.store
```

Use the following keytool command to change the server certificate private key entry password:

```
keytool -keypasswd -alias tomcat -keystore <installation_
folder>/central/var/security/key.store
```

Use the following keytool command to change the TrustStore password:

```
keytool -storepasswd -keystore <installation_
folder>/central/var/security/client.truststore
```

4. Change the passwords also in the **server.xml** file located in **<installation dir>/central/tomcat/conf/server.xml**.

- a. Locate the HTTPS connector. For example:

```
keyPass="changeit" keystoreFile="C:/Program Files/Hewlett-Packard/HP
Operations Orchestration/central/var/security/key.store"
keystorePass="changeit" keystoreType="JKS" maxThreads="200" port="8443"
protocol="org.apache.coyote.http11.Http11NioProtocol" scheme="https"
secure="true" sslProtocol="TLSv1.2"
sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2" truststoreFile="C:/Program
Files/Hewlett-Packard/HP Operations
Orchestration/central/var/security/client.truststore"
truststorePass="changeit" truststoreType="JKS"/>
```

Change the required password.

- **keyPass** - the password used to access the server certificate private key from the specified key.store file. The default value is "changeit".
- **keystorePass** - the password used to access the specified key.store file. The default value is the value of the **keyPass** attribute.

Note: It is recommend not to use the same password as the **keyPass**, and to use a strong password.

- **truststorePass** - the password to access the TrustStore (that includes all of the trusted CAs). The default is the value of the **javax.net.ssl.trustStorePassword** system property. If that property is null, no TrustStore password will be configured. If an invalid TrustStore password is specified, a warning will be logged and an attempt will be made to access the TrustStore without a password, which will skip validation of the TrustStore contents.

- b. Save the file.

5. Edit the **central-wrapper.conf** file located in **<installation dir> central\conf\central** and replace the password of the TrustStore with new password in encrypted or obfuscated form. Examples:

```
wrapper.java.additional.<x>=-Djavax.net.ssl.trustStorePassword={ENCRYPTED}
<encrypted_password>
```

```
wrapper.java.additional.<x>=-Djavax.net.ssl.trustStorePassword={OBFUSCATED}  
<obfuscated_password>
```

For information about how to encrypt or obfuscate a password, see ["Encrypting and Obfuscating Passwords"](#) below.

6. Start the Central service.

Changing the RAS, OOSH, and Studio TrustStore Passwords

Note: You should change the KeyStore, TrustStore, and server certificate password using Keytool, before completing the following steps.

- **To change the standalone RAS TrustStore password:** Edit the `ras-wrapper.conf` file and change the password of the TrustStore.
- **To change the OOSH TrustStore password:** Edit the `oosh.bat` file and change the password of the TrustStore.
- **To change the Studio TrustStore password:** Add the property `client.truststore.password` with the password in obfuscated format to the `Studio.properties` file in the `<user>/.` folder.

```
client.truststore.password={OBFUSCATED}6L9+NqBjKYp5heuvMEzg0g==
```

If this property is not defined, Studio will fall back to the system property `javax.net.ssl.trustStorePassword` for the TrustStore password.

For information about how to obfuscate a password, see ["Encrypting and Obfuscating Passwords"](#) below.

Encrypting and Obfuscating Passwords

You can encrypt or obfuscate a password using the `encrypt-password` script, which is located in `<installation_folder>/central/bin`.

Our recommendation is to use encryption.

Important! After using the `encrypt-password` script, clear the command history.

This is because, on a Linux OS, the password parameter will be stored in cleartext under `/$USER/.bash_history` and accessible by the `history` command.

Encrypting Passwords

1. Locate the `encrypt-password` script, in `<installation_folder>/central/bin`.
2. Run the script with the `-e -p <password>` option, where `password` is the password you want to encrypt.

Note: You can either use `-p` as the flag to encrypt the password or `--password`.

The encrypted password should appear as follows:

```
{ENCRYPTED}<some_chars>.
```

Obfuscating Passwords

1. Locate the `encrypt-password` script, in `<installation_folder>/central/bin`.
2. Run the script with the `-o <password>` option, where `password` is the password that you want to obfuscate.

The obfuscated password should appear as follows:

```
{OBFUSCATED}<some_chars>.
```

Creating a Prompt for the Password

It is recommended to run the `encrypt-password` script without providing the `-p` argument. For example:

```
C:\Program Files\Hewlett-Packard\HP Operations Orchestration\central\bin>encrypt-password
Password (typing will be hidden):
Confirm password (typing will be hidden):
{ENCRYPTED}gAkPCLQsYDhoR1Y2q9BjCQ==
C:\Program Files\Hewlett-Packard\HP Operations Orchestration\central\bin>
```

This will create a prompt for the hidden password inputs.

Removing Vulnerable Ciphers from the SSL supported Ciphers

The DES and Triple DES ciphers, as used in the TLS protocol have a birthday bound of approximately four billion blocks, which makes it easier for remote attackers to obtain cleartext data via a birthday attack against a long-duration encrypted session, as demonstrated by an HTTPS session using Triple DES in CBC mode, aka a "Sweet32" attack.

For more information about this attack, see <https://sweet32.info/>.

To disable RC4, DES and Triple DES ciphers in OO:

1. Open the `$JRE_HOME/lib/security/java.security` file.
2. Remove the comments and change the parameters according to the example below:

```
jdk.certpath.disabledAlgorithms=DES, DESede, RC4, MD2, RSA keySize < 1024
jdk.tls.disabledAlgorithms=DES, DESede, RC4, MD5, DSA, RSA keySize < 1024
```

3. Restart the OO Central server.

For more information, see <http://stackoverflow.com/questions/18589761/restrict-cipher-suites-on-jrelevel>.

After upgrading from an earlier version of OO 10.x, repeat these steps.

Changing the HTTP/HTTPS Ports or Disabling the HTTP Port

The file `server.xml` under `[OO_HOME]/central/tomcat/conf` contains two elements named `<Connector>` under the element `<Service>`. These connectors define or enable the ports that the server are listening to.

Each connector configuration is defined through its attributes. The first connector defines a regular HTTP connector and the second defines an HTTPS connector.

By default, the connectors look as follows.

HTTP connector:

```
<Connector URIEncoding="UTF-8" compression="on" connectionTimeout="20000"
port="8080" protocol="org.apache.coyote.http11.Http11NioProtocol"
redirectPort="8443"/>
```

HTTPS connector:

```
<Connector SSLEnabled="true" URIEncoding="UTF-8" clientAuth="false"
```

```
compression="on" keyAlias="tomcat" keyPass="changeit" keystoreFile="C:/Program
Files/Hewlett-Packard/HP Operations
Orchestration/central/var/security/key.store" keystorePass="changeit"
keystoreType="JKS" maxThreads="200" port="8443"
protocol="org.apache.coyote.http11.Http11NioProtocol" scheme="https"
secure="true" sslProtocol="TLSv1.2" sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2"
truststoreFile="C:/Program Files/Hewlett-Packard/HP Operations
Orchestration/central/var/security/client.truststore" truststorePass="changeit"
truststoreType="JKS"/>
```

By default, both are enabled.

Important! If you change or disable one of the Central ports in the **server.xml** file, you will also need to update the **central-wrapper.conf** file and each **RAS-wrapper.conf** file to point to the Central URL with the updated port. Otherwise, all your flows will fail when run from Central. Also, make sure to check the load balancer configurations.

Changing Port Values

To change the values of one of the ports:

1. Edit the **server.xml** file located in **<installation_dir>/central/tomcat/conf/server.xml**.
2. Locate the HTTP or HTTPS connector, and adjust the **port** value in the line.

Note: If you are keeping both HTTP and HTTPS active and you want to change the HTTPS port, you will need to change the **redirectPort** value for the HTTP connector and the **port** value for the HTTPS connector.

3. Save the file.
4. Restart Central.

Disabling the HTTP Port

You might want to disable the HTTP port, for security reasons, so that the only communication channel will be on TLS and encrypted.

1. Edit the **server.xml** file located in **<installation_dir>/central/tomcat/conf/server.xml**.
2. Locate the HTTP connector, and delete or comment out the line.
3. Import the trusted root certificate authority (CA) to the Central **client.truststore** file, if it doesn't already exist in the CA list:

```
keytool -importcert -alias <any_alias> -keystore <path to the  
client.truststore> -file <certificate_name.cer> -storepass <changeit>
```

Note: If you are using a well known root CA (like Verisign) you do not have to perform this step, because the certificate will already be in the `client.truststore` file.

4. Save the file.
5. Restart Central.

Note: It is also possible to disable the HTTP port during installation.

Troubleshooting the HTTPS Connector

If the server doesn't start, open the `wrapper.log` file and look for an error in `ProtocolHandler` [`"http-nio-8443"`].

This can happen when Tomcat is initializing or starting the connector. There are many variations but the error message can provide information.

All the HTTPS connector parameters are in the Tomcat configuration file located at `C:\HPE\oo\central\tomcat\conf\server.xml`.

Open the file and scroll to the end, until you see the HTTPS connector:

```
<Connector SSLEnabled="true" clientAuth="false" keyAlias="tomcat"  
keystoreFile="C:/HPE/oo/central/var/security/keystore.p12" keystorePass="tomcat-  
keystore-password" keystoreType="PKCS12" maxThreads="200" port="8443"  
protocol="org.apache.coyote.http11.Http11NioProtocol" scheme="https" secure="true"  
sslProtocol="TLSv1.2" sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2"/>
```

See if there is any mismatch in the parameters, by comparing them to the parameters you entered in the previous steps.

Client Certificate Authentication (Mutual Authentication)

The most common use of X.509 certificate authentication is in verifying the identity of a server when using TLS, most commonly when using HTTPS from a browser. The browser automatically checks that the certificate presented by a server has been issued by one of a list of trusted certificate authorities, which it maintains.

You can also use TLS with mutual authentication. The server requests a valid certificate from the client as part of the TLS handshake. The server authenticates the client by checking that its certificate is

signed by an acceptable authority. If a valid certificate has been provided, it can be obtained through the servlet API in an application.

Configuring Client Certificate Authentication in Central

Before you configure the client certificate authentication in Central, make sure you have configured the TLS server certificate, as described in ["Working with Server and Client Certificates" on page 109](#).

By default, the user certificate is matched against internal users of Operations Orchestration (OO) Central. If you want to match the certificate against LDAP users, make sure you set the desired LDAP as default LDAP.

Set the `clientAuth` attribute to `true` if you want the TLS stack to request a valid certificate chain from the client before accepting a connection.

Set the `clientAuth` attribute to `want` if you want the TLS stack to request a client certificate. When the attribute is set to `want` and:

- The certificate is not presented, then the authentication does not fail and the user is redirected to the login page.
- An invalid certificate is presented, then the authentication fails and the user is not redirected to the login page.

A `false` value (default) does not require a certificate chain unless the client requests a resource protected by a security constraint that uses CLIENT-CERT authentication. (For more information, see the Apache Tomcat Configuration Reference).

Set the **Certificate Revocation List (CRL)** file. This can contain several CRLs. In the operation of some cryptosystems, usually public key infrastructures (PKIs), a certificate revocation list (CRL) is a list of certificates (or more specifically, a list of serial numbers for certificates) that have been revoked, and therefore, entities presenting those (revoked) certificates should no longer be trusted.

Note: The following procedure uses the Keytool utility that is located in `<installation dir>/java/bin/keytool`.

1. Stop the Central server.
2. Import the appropriate root certificate (CA) into Central `client.truststore`: `<installation dir>/central/var/security/client.truststore`, if it doesn't already exist in the CA list (by default, all the well known CAs are there). For example:

```
keytool -importcert -alias <any_alias> -keystore <path>/client.truststore -file  
<certificate_path> -storepass <changeit>
```

3. Edit the **server.xml** file located in **<installation_dir>/central/tomcat/conf/server.xml**.
4. Set the **clientAuth** attribute in the Connector tag to **want** or **true**. The default is **false**.

For example:

```
<Connector SSLEnabled="true" URIEncoding="UTF-8" clientAuth="false"  
compression="on" keyAlias="tomcat" keyPass="changeit"  
keystoreFile="C:/Program Files/Hewlett-Packard/HP Operations  
Orchestration/central/var/security/key.store" keystorePass="changeit"  
keystoreType="JKS" maxThreads="200" port="8443"  
protocol="org.apache.coyote.http11.Http11NioProtocol" scheme="https"  
secure="true" server="00" sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2"  
sslProtocol="TLSv1.2" truststoreFile="C:/Program Files/Hewlett-Packard/HP  
Operations Orchestration/central/var/security/client.truststore"  
truststorePass="changeit" truststoreType="JKS"/>
```

Note:

- We recommend starting the server at the end of this procedure, but note that it is also possible to start the server at this point.
 - The OO 9x backwards compatible SOAP/REST APIs are not supported in cases where Client authentication is required.
5. (Optional) Add the **crlFile** attribute to define the certificate revocation list file for the TLS certificate validation, for example:

```
crlFile="<path>/crlname.<crl/pem>"
```

The file can be with the **.crl** extension for a single certificate revocation list or with the **.pem** (PEM CRL format) extension for one or more certificate revocation lists. The PEM CRL format uses the following header and footer lines:

```
-----BEGIN X509 CRL-----  
-----END X509 CRL-----
```

Example of the **.pem** file structure for one CRL (for more than one, concatenate another CRL block):

```
-----BEGIN X509 CRL-----  
MIIBbzCB2QIBATANBgqhkiG9w0BAQUFADBeMQswCQYDVQQGEwJVUzEYMBYGA1UE  
ChMPVS5TLiBhb3Zlcm5tZW50MQwwCgYDVQQLEwNEb0QxEDA0BgNVBAsTB1Rlc3Rp
```

```
bmcxFTATBgNVBAMTDFRydXN0IEFuY2hvchcNOTkwMTAxMTIwMTAwWhcNNDgwMTAx  
MTIwMTAwWjAiMCAcAScXDTk5MDEwMTEyMDAwMFowDDAKBgNVHRUEAwoBAaAjMCEw  
CgYDVLR0UBAMCAQEWYDVR0jBAwwCoAIq5rr+cLnVI8wDQYJKoZIhvcNAQEFBQAD  
gYEAC71qZwejJRW7QvzH11/7cYcL3racgMxH3PSU/ufvyLk7ahR++RtHary/WeCv  
RdyznLiIOA8ZBiguWtVPqsNysNn7WLoFQIVa+/TD3T+1ece4e1NwGQvj5Q+e2wRt  
GXg+gCuTjTKUFfKRnWz707RyiJKKIm0jtAF4RkCpLebNChY=  
-----END X509 CRL-----
```

6. Edit the **central-wrapper.conf** file, located in **<installation dir> central\conf\central**.

Uncomment the following properties and set the client certificate location and password to a client certificate with an administrator user.

```
#wrapper.java.additional.23=-Djavax.net.ssl.keyStore="%CENTRAL_  
HOME%/var/security/certificate.p12"  
  
#wrapper.java.additional.24.stripquotes=TRUE  
  
#wrapper.java.additional.25=-Djavax.net.ssl.keyStorePassword={OBFUSCATED}  
ZUoMreNLw6qI0yzX7g5YKw==  
  
#wrapper.java.additional.26=-Djavax.net.ssl.keyStoreType=PKCS12
```

For information about how to encrypt or obfuscate a password, see ["Encrypting and Obfuscating Passwords" on page 119](#).

7. Start the Central server.

Note: For each client certificate, you need to define a user, either an internal user or LDAP user. The name of the user should be defined in the certificate attributes. The default is value of the CN attribute. See the [Processing a Certificate Principal](#) section for more details.

Note that even if OO is set up with multiple LDAP configurations, it is only possible to authenticate the user using the client certificate attributes with the default LDAP.

Updating the Configuration of a Client Certificate in RAS

The client certificate is configured during the installation of the RAS. However, if you need to update the client certificate, you can do this manually in the **ras-wrapper.conf** file.

Prerequisite: You must import the CA root certificate of Central into the RAS TrustStore. See ["Importing a CA Root Certificate to a RAS TrustStore" on page 112](#).

To update the configuration of the client certificate in an external RAS:

1. Stop the RAS server.
2. Open the **ras-wrapper.conf** file from **<installation dir>ras/conf/ras-wrapper.conf**.

3. Change the following according to your client certificate:

```
wrapper.java.additional.<x>=-Djavax.net.ssl.keyStore=<installation  
dir>/var/security/certificate.p12"
```

```
wrapper.java.additional.<x>=-Djavax.net.ssl.keyStorePassword={OBFUSCATED}  
<obfuscated_password>
```

```
wrapper.java.additional.<x>=-Djavax.net.ssl.keyStoreType=PKCS12
```

4. Start the RAS server.

Important Notes! The X.509 client certificate needs to have the principal name of the RAS, which is the RAS ID (see [Processing a Certificate Principal](#)).

You can find the RAS ID under the **Topology** tab in Central. See "Setting Up Topology – Workers" in the *_HPc_Basic_Variables._HP_Product_AcronymCentral User Guide*.

In OO 10.20 and later, the `keyStorePassword` parameter is obfuscated by default if the password was kept as default. You can change this parameter and store it in either clear text or obfuscated. See "[Encrypting and Obfuscating Passwords](#)" on page 119.

Configuring a Client Certificate in Studio Remote Debugger

Prerequisite: You must import the CA root certificate of Central into the Studio Debugger TrustStore. See "[Importing a CA Root Certificate to the Studio TrustStore](#)" on page 114.

To configure the client certificate in the Studio Remote Debugger:

1. Close Studio.
2. Edit the **Studio.I4j.ini** file from **<installation dir>/studio**.
3. Change the following according to your client certificate:

```
-Djavax.net.ssl.keyStore="<installation  
dir>/studio/var/security/certificate.p12"
```

```
-Djavax.net.ssl.keyStorePassword={OBFUSCATED}<obfuscated_password>
```

```
-Djavax.net.ssl.keyStoreType=PKCS12
```

4. Start Studio.

Notes:

- In OO 10.20 and later, the `keyStorePassword` parameter is obfuscated by default if the password was kept as default. You can change this parameter and store it in either clear text or

obfuscated. See ["Encrypting and Obfuscating Passwords" on page 119](#).

- For the client certificate, you need to define a user, either an internal user or LDAP user. The name of the user should be defined in the certificate attributes. The default is value of the CN attribute. See the [Processing a Certificate Principal](#) section for more details.
- Note that even if OO is set up with multiple LDAP configurations, it is only possible to authenticate the user using the client certificate attributes with the default LDAP. Central will first try to authenticate the user with the default LDAP, and if this fails, will try to authenticate within the OO internal domain.

Configuring a Client Certificate in OOSH

Prerequisite: You must import the CA root certificate of Central into the OOSH TrustStore. See ["Importing a CA Root Certificate to the OOSH TrustStore" on page 113](#).

1. Stop OOSH.
2. Edit the `oosh.bat` from `<installation dir>/central/bin` (in a standalone OOSH, it is under the main OOSH directory).
3. Change the following according to your client certificate:

```
-Djavax.net.ssl.keyStore="<installation dir>/var/security/certificate.p12"  
-Djavax.net.ssl.keyStorePassword={OBFUSCATED}<obfuscated_password>  
-Djavax.net.ssl.keyStoreType=PKCS12
```

4. Start OOSH.

Note:

In OO 10.20 and later, the `keyStorePassword` parameter is obfuscated by default if the password was kept as default. You can change this parameter and store it in either clear text or obfuscated. See ["Encrypting and Obfuscating Passwords" on page 119](#).

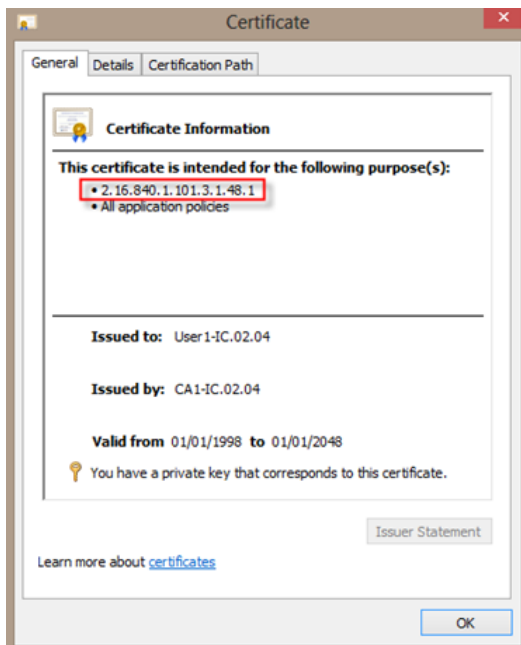
For the client certificate, you need to define a user, either an internal user or LDAP user. The name of the user should be defined in the certificate attributes. The default is value of the CN attribute. See the [Processing a Certificate Principal](#) section for more details.

Note that even if OO is set up with multiple LDAP configurations, it is only possible to authenticate the user using the client certificate attributes with the default LDAP. Central will first try to authenticate the user with the default LDAP, and if this fails, will try to authenticate within the OO internal domain.

Processing Certificate Policies

OO handles the processing of certificate polices for the end certificate.

- You can set the purpose string in the certificate.
- OO lets you add the policy string(s) as a configuration item and check the policy string of each end certificate. If it does not match, reject the certificate.
- Enable or disable the certificate policy verification by adding the following configuration item:
`x509.certificate.policy.enabled=true/false` (default is false).
- Define the policy list by adding the following configuration item:
`x509.certificate.policy.list=<comma_separated_list>` (the default is an empty list).



For more information about how to change OOsystm properties, see the *OO Shell Guide*.

Processing a Certificate Principal

You can define how to get the principal from a certificate using a regular expression match against the Subject. The regular expression should contain a single group. The default expression `CN=(.?)` matches the common name field. For example, `CN=Jimi Hendrix, OU=` assigns a user name of Jimi Hendrix.

- The matches are case-insensitive.
- The principal of the certificate is the user name in OO (LDAP or internal user).
- To change the regular expression, change the configuration item:
`x509.subject.principal.regex`.

Enabling `_HPc_Basic_Variables._HP_Product_Acronym` to Read from the Subject Alternative Name Field in a Certificate

You can enable `_HPc_Basic_Variables._HP_Product_Acronym` to read from the Subject Alternative Name field in a certificate, using the configuration item `x509.principal.lookup.field`.

This configuration item controls which certificate field is used to extract the username.

Possible values are:

- `subjectDN` - represents the Subject field of the certificate, meaning that `_HPc_Basic_Variables._HP_Product_Acronym` keeps its default behavior and attempts to extract the username from the **Subject** field. This is the default value.
- `subjectAltNames.otherName.principalName` - represents the User Principal Name (OID 1.3.6.1.4.1.311.20.2.3) contained in the Other Name entry of the Subject Alternative Names certificate extension. For CAC Authentication, it might be required to use the value of the User Principal Name, so you would use this value.

For more information about how to change OO configuration items, see the *OO Shell (OOSH) User Guide*.

Signing OO Content Packs

Code signing is a mechanism whereby publishers of software and content use a certificate-based digital signature to verify their identities to users of the code, thus allowing users to decide whether or not to install it based on whether they trust the publisher.

It is important that Operations Orchestration content packs are code signed with your digital signature for the following security benefits:

- Operations Orchestration administrators can identify your company as the author, publisher or distributor of the content pack, at the time of deploying content packs to Operations Orchestration Central.

- It provides integrity of the content pack, that is it protects the content packs from being tampered with (if the content pack is changed, the digital signature is invalidated).

A digital certificate is an electronic document issued by a Certificate Authority (CA).

It contains the public key for a digital signature and specifies the identity associated with the key, such as the name of an organization.

The certificate is used to confirm that the public key belongs to the specific organization. The CA acts as the guarantor.

Digital certificates must be issued by a trusted authority and are only valid for a specified time. They are required in order to create a digital signature.

A typical option to store the digital signature file is inside a PKCS#12 file that has .pfx extension, for example **digital_signature.pfx**.

What do you want to do?

Sign content packs using the jarsigner utility

1. Download Java SE 8u102 from here:
<http://www.oracle.com/technetwork/java/javase/downloads/index.html>
2. Install it to a suitable location on your local file system and locate your jarsigner utility at the following location <JDK_INSTALLTION>/bin.
3. Use the digital signature to sign the content pack file .jar.

Convert the digital signature file from PKCS#12 format to JKS format used by the jarsigner utility

1. Using the keytool utility located in <oo_installation>/java/bin/keytool, run the following command:

```
keytool -importkeystore -srckeystore  
<digital_signature.pfx> -srcstoretype pkcs12 -srcalias  
<certificate_alias_pfx> -destkeystore  
<digital_signature.jks> -deststoretype jks -deststorepass  
<password_jks> -destalias <certificate_alias_jks>
```

where:

Name	Description
certificate_alias_pfx	The alias inside the source digital file in pkcs12 format identifying the private key to be used to sign the content pack JAR file, and the key's associated certificate
digital_signature.jks	The name of the output digital file in JKS format
deststorepass	The password for the output digital file in JKS format
certificate_alias_jks	The alias inside the output digital file in JKS format identifying the private key to be used to sign the content pack JAR file, and the key's associated certificate

2. Then, run the following command:

```
jarsigner -keystore <digital_signature.jks>  
-signedjar <signed-content-pack.jar> <content_pack.jar>  
<certificate_alias_jks>
```

where:

Name	Description
digital_signature.jks	The name of the output digital file in JKS format
signed-content-pack.jar	The output signed content pack file JAR
content_pack.jar	The JAR file to be used as input for signing
certificate_alias_jks	The alias inside the output digital file in JKS format identifying the private key to be used to sign the content pack JAR file, and the key's associated certificate.

Verifying the Integrity of the Studio Tools JARs

There are JAR files located under the <OO_installation_folder>/studio/tools/lib folder that provide the executable code for OOSHA, wizards and hpln-index-generator tools.

In order to verify the integrity of the JAR files, to ensure that the JARs have not been tampered with since the Operations Orchestration installation, you can verify their digital signature.

A digital certificate is an electronic document issued by a Certificate Authority (CA).

It contains the public key for a digital signature and specifies the identity associated with the key, such as the name of an organization.

The certificate is used to confirm that the public key belongs to the specific organization. The CA acts as the guarantor.

Digital certificates must be issued by a trusted authority and are only valid for a specified time. They are required in order to create a digital signature.

See the *Operations Orchestration Studio Wizards Guide* for more information.

What do you want to do?

Verify the digital signature of a Studio tools JAR

1. Download Java SE 8u102 from here:

<http://www.oracle.com/technetwork/java/javase/downloads/index.html>

2. Install it in a suitable location on your local file system and locate your jarsigner utility under:

<JDK_INSTALLATION>/bin

3. Run the following command from the command line

```
jarsigner -verify -keystore  
<oo_installation>/central/var/security/internal.truststore  
-strict -verbose -certs <jar_file>
```

where <jar_file> is any one of the following jars:

- o hpln-index-generator.jar
- o oosha-jar-with-dependencies.jar
- o ps-wizard-jar-with-dependencies.jar
- o rest-wizard-jar-with-dependencies.jar
- o shell-wizard-jar-with-dependencies.jar
- o third-party-cp-wizard-jar-with-dependencies.jar
- o ws-wizard-jar-with-dependencies.jar

The output of the command is expected to end in "jar verified". Additionally, all JAR file entries listed in the output must be accompanied by the "smk" label.

Configuring OO for FIPS 140-2 Level 1 Compliance

The section below explains how to configure Operations Orchestration (OO) to be compliant with Federal Information Processing Standards (FIPS) 140-2 Level 1.

FIPS 140-2 is a standard for security requirements for cryptographic modules defined by the National Institute of Standards Technology (NIST). To view the publication for this standard, go to: csrc.nist.gov/publications/fips/fips140-2/fips1402.pdf.

After you have configured OO for FIPS 140-2 compliance, OO uses the following security algorithm:

- Symmetric-key algorithm: AES256
- Hashing algorithm: SHA256

OO uses the security provider: RSA BSAFE Crypto software version 6.2.1. This is the only supported security provider for FIPS 140-2.

Note: Once you have configured OO to be compliant with FIPS 140-2, you can only revert back to the standard configuration if you re-install OO.

Prerequisites

Note for upgraders:

If you are upgrading from an installation of OO 10.10 (and later) that was already configured with FIPS, see [Prerequisite Steps for Upgraders](#).

Before configuring OO to be compliant with FIPS 140-2, perform the following steps:

Note: In order to be FIPS140-2 compliant, you need to turn off LWSSO.

1. Verify that you are configuring a new installation of OO version 10.10 or later to be compliant with FIPS 140-2, and that it is not in use.

You cannot configure an installation of OO that is in use (whether version 9.x or 10.x).

2. Verify that when OO was installed, it was configured not to start the Central server after installation:
 - In a silent installation, the `should.start.central` parameter was set to **no**.
 - In a wizard installation, in the **Connectivity** step, the **Do not start Central server after**

installation check box was selected.

Connectivity

Configure the Central Server port numbers and SSL properties

HTTP 8080

HTTPS 8443

Provide a secure SSL certificate (when not provided, a self-signed certificate is used)

Secure keystore Browse...

The secure keystore should be in PKCS12 format and include both certificate and private key. Usually this is a file with a .pfx or .p12 extension. Consult your Certificate Authority for more details.

Keystore password

Do not start Central server after installation (Must be checked when you want to configure HP OO to be compliant with FIPS 140-2.)

3. Back up the following directories:
 - o **<installation dir>\central\tomcat\webapps\oo.war**
 - o **<installation dir>\central\tomcat\webapps\PAS.war**
 - o **<installation dir>\central\conf**
 - o **<installation dir>\java** (the entire **java** folder should be backed-up)
4. Download **Server Oracle JRE 8** from <http://www.oracle.com/technetwork/java/javase/downloads/server-jre8-downloads-2133154.html>, and replace the **OpenJDK (Zulu) JRE** with the **Server Oracle JRE**.
 - a. Delete everything inside the **<installation dir>\java** folder.
 - b. Extract the downloaded archive.
 - c. Copy the **JRE** folder content to **<installation dir>\java**.
5. Download and install the Java Cryptographic Extension (JCE) Unlimited Strength Jurisdiction Policy Files from the following site:
<http://www.oracle.com/technetwork/java/javase/downloads/server-jre8-downloads-2133154.html>

Note: See the **ReadMe.txt** file from the downloaded content for information on how to deploy the files and upgrade the JRE used by OO.
6. Install the RSA BSAFE Crypto software files. On the system on which OO is installed, copy the following to **<oo_jre>\lib\ext** (where **<oo_jre>** is the directory in which the JRE used by OO is installed. By default, this is **<installation dir>\java**).
 - o **<installation dir>\central\lib\cryptojce-6.2.1.jar**
 - o **<installation dir>\central\lib\cryptojcommon-6.2.1.jar**
 - o **<installation dir>\central\lib\jcmFIPS-6.2.1.jar**

Prerequisite Steps for Upgraders

1. Download the Server Oracle JRE 8 and replace the OpenJDK (Zulu) JRE with the Server Oracle JRE.
 - a. Delete everything inside the **<upgrade dir>\JAVA** folder.
 - b. Extract the downloaded archive.
 - c. Copy the **JRE** folder content to **<upgrade dir>\JAVA**.

<http://www.oracle.com/technetwork/java/javase/downloads/server-jre8-downloads-2133154.html>

2. Download and install the Java Cryptographic Extension (JCE) Unlimited Strength Jurisdiction Policy Files from the following site:

<http://www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html>

See the **ReadMe.txt** file from the downloaded content for information on how to deploy the files and upgrade the JRE used by OO.

3. Install the RSA BSAFE Crypto software files. On the system on which OO is installed, copy the following files to **<oo_jre>\lib\ext**:

(where **<oo_jre>** is the directory in which the JRE that is used by the OO upgrader. By default, this is **<upgrade dir>\java**).

- **<installation dir>\central\lib\cryptojce-6.2.1.jar**
- **<installation dir>\central\lib\cryptojcommon-6.2.1.jar**
- **<installation dir>\central\lib\jcmFIPS-6.2.1.jar**

Next, follow the steps in the "Configure the Properties in the Java Security File" section in "[Configuring Compliance with FIPS 140-2](#)" below.

Configuring Compliance with FIPS 140-2

The following list shows the procedures that you need to perform in order to configure OO) to be compliant with FIPS 140-2:

1. [Configure the Properties in the Java Security File.](#)
2. [Configure the encryption.properties File and Enable FIPS Mode.](#)
3. [Create FIPS-CompliantOO Encryption.](#)

4. [Re-encrypt the database password with the new encryption.](#)
5. [Start OO.](#)

Step 1: Configure the Properties in the Java Security File

Edit the Java security file for the JRE to add additional security providers and configure the properties for FIPS 140-2 compliance.

Note: The upgrade to OO 10.x completely replaces the installed JRE files. Therefore, if you are upgrading to 10.x, you must complete the following steps.

Note: If you are upgrading from an installation of OO 10.10 and later that was already configured with FIPS, you must follow the "Prerequisite Steps for Upgraders" section in "[Configuring OO for FIPS 140-2 Level 1 Compliance](#)" on page 134, and then follow the steps here, where `<oo_jre>` is the JRE included in the upgrade (in the location `<upgrade_dir>\JAVA`).

Make sure to make all the changes in the `java` folder inside the extracted `upgrade` folder.

Open the `<oo_jre>\lib\security\java.security` file in an editor and perform the following steps:

1. For every provider listed, in the format `security.provider.<nn>=<provider_name>`, increment the preference order number `<nn>` by two.

For example, change a provider entry from:

```
security.provider.1=sun.security.provider.Sun
```

to

```
security.provider.3=sun.security.provider.Sun
```

2. Add a new default provider (RSA JCE). Add the following provider at the top of the provider list:

```
security.provider.1=com.rsa.jsafe.provider.JsafeJCE
```

3. Add the RSA BSAFE SSL-J Java Secure Sockets Extension (JSSE) Provider:

```
security.provider.2=com.rsa.jsse.JsseProvider
```

4. Copy and paste the following line into the `java.security` file to ensure **RSA BSAFE** is used in FIPS 140-2 compliant mode:

```
com.rsa.cryptoj.fips140initialmode=FIPS140_SSL_MODE
```

You can paste this line anywhere in the `java.security` file.

5. Because the default DRBG algorithm ECDRBG128 is not safe (according to NIST), set the

security property **com.rsa.crypto.default** to **HMACDRBG**, by copying the following line into the **java.security** file:

```
com.rsa.crypto.default.random=HMACDRBG
```

You can paste this line anywhere in the **java.security** file.

6. Save and exit the **java.security** file.

Step 2: Configure the encryption.properties File and Enable FIPS Mode

The OO encryption properties file must be updated to be FIPS 140-2 compliant.

1. Back up the **encryption.properties** file, which is located in **<installation dir>\central\var\security**.
2. Open the **encryption.properties** file in a text editor. For example, edit the following file:

```
C:\Program Files\Hewlett-Packard\HP Operations  
Orchestration\central\var\security\encryption.properties.
```

3. Locate **keySize=128** and replace it with **keySize=256**.
4. Locate **secureHashAlgorithm=SHA1** and replace it with **secureHashAlgorithm=SHA256**.
5. Locate **FIPS140ModeEnabled=false** and replace it with **FIPS140ModeEnabled=true**.

Note: If **FIPS140ModeEnabled=false** does not exist, add **FIPS140ModeEnabled=true** as a new line to the end of the file.

6. Save and close the file.

Step 3: Create FIPS-Compliant Encryption

To create or replace the OO encryption store file, so that it is FIPS-compliant, see ["Replacing the FIPS Encryption" on the next page](#).

Note: AES has three approved key lengths: 128/192/256 by NIST SP800-131A publication.

The following secure hash algorithms are supported in FIPS: SHA1, SHA256, SHA384, SHA512.

Note: It is recommended to change the passwords of the **key.store** (and its private key entry) and the truststore. See ["Changing and Encrypting/Obfuscating the KeyStore/TrustStore Password" on page 117](#)

Note: It is recommended to delete all the default CA root certificates that are not in use from the OO truststore. (The **client.truststore** is located at **<installation>/central/var/security**.)

Note: If you work with Client Certificate, the certificate should be generated with the FIPS-compliant RSA JCE provider, and with the secure hash algorithms that are supported in FIPS, as listed above.

Step 4: Re-encrypt the database password with the new encryption

Re-encrypt the database password, as described in the *OO Administration Guide*, in "Changing the Database Password Guide".

Step 5: Start OO

Replacing the FIPS Encryption

OO Central and RAS comply with Federal Information Processing Standard 140-2 (FIPS 140-2), which defines the technical requirements to be used by federal agencies when these organizations specify cryptographic-based security systems for protection of sensitive or valuable data.

After a fresh installation of OO, you have the option to change the FIPS encryption key.

Note: This procedure is only for fresh installations. You cannot perform it after an upgrade.

Changing the FIPS Encryption Key on Central

Use the **generate-keys.bat/sh** file to replace the FIPS encryption key in the encryption repository.

Note: This process backs up the **encryption_repository** file, so you must have the relevant write permissions.

1. Go to **<Central installation folder>/var/security**.
2. Back up the **encryption_repository** file and delete it from the **<Central installation folder>/var/security** folder.
3. Go to **<Central installation folder>/bin**.
4. Run the **generate-keys** script.
5. Press the **Y** key to proceed.

A new master key is generated in **<Central installation folder>/var/security/encryption_repository**.

Note: If you prefer to run the **generate-keys** script without pausing for the user to type **Y** or **N**, use the silent mode flag **-s** when running the script.

Changing the RAS Encryption Properties

If the installation of the RAS is in a new location, you need to complete all the steps below.

Note: These changes are only valid if you working on a new RAS installation after you have changed the Central encryption properties.

To change the RAS encryption properties:

1. Complete all the steps in the "Prerequisites" section in ["Configuring OO for FIPS 140-2 Level 1 Compliance" on page 134](#).
2. Complete all the steps in the "Configure the Properties in the Java Security File" in ["Configuring Compliance with FIPS 140-2" on page 136](#).
3. Copy the current **encryption.properties** file from **<installation dir>\ras\var\security** to the **<installation dir>\ras\bin** folder.
4. Using any text editor, edit and change the **encryption.properties** file as required.

For more information, see "Configure the encryption.properties File and Enable FIPS Mode" in ["Configuring Compliance with FIPS 140-2" on page 136](#).

5. Save the changes.
6. Open a command line prompt in the folder **<installation dir>\ras\bin**.
7. Run **oosh.bat**.
8. Run the OOShell command: `replace-encryption --file encryption.properties`

Note: If you copied the **encryption.properties** file to a different folder, make sure you enter the correct location in the OOShell command.

9. Restart the RAS service.

Configuring the TLS Protocol

You can configure OO to define the supported TLS protocol version. By default, OO allows TLS v1, TLS v1.1 and TLS v1.2, but you can narrow this down.

Note: SSLv3 and other versions of SSL are not supported.

1. Open the `<installation_folder>/central/tomcat/conf/server.xml` file .
2. Locate the SSL connector (at the end of the file).
3. Edit the default value of `sslEnabledProtocols`. For example, change
`sslEnabledProtocols="TLSv1,TLSv1.1,TLSv1.2"` to
`sslEnabledProtocols="TLSv1.2"`
4. Restart the server.

Preventing Flows from Accessing the Central/RAS Local File System

You should modify the wrapper configuration and `java.policy` files of Central or RAS, in order to prevent flows from accessing the local file system of Central or RAS, and gaining access to sensitive resources.

Note: In order to exploit this scenario, a user would need to have both deployment and triggering permissions, in addition to entitlement on flows or the ability to entitle flows. Users with such permissions are likely to be trustworthy users.

To protect from this scenario:

1. In the wrapper configuration file of Central or RAS (`<installation_folder>/<ras/central>/conf/<central/ras>-wrapper.conf`), add the `wrapper.java.additional.<nn>` parameter as follows:
`wrapper.java.additional.<nn>=-Djava.security.manager`
Replace `<nn>` with the number after the last number.
2. In the `java.policy` file (located at `<installation_folder>/java/lib/security/java.policy`), add the following. This allows access to the minimum resources that are required by Operations Orchestration and prevents access to the local file system of Central/RAS that contains sensitive data.

```
grant codebase "file:${oo.home}/bin/-" {
    permission java.security.AllPermission;
};
grant codebase "file:${oo.home}/lib/-" {
    permission java.security.AllPermission;
};
```

```
grant codebase "file:${oo.home}/tomcat/-" {
    permission java.security.AllPermission;
};

grant codebase "file:${oo.home}/var/cache/-" {
    permission java.lang.RuntimePermission "getClassLoader";
    permission java.io.FilePermission "${oo.home}/var/cache/-",
        "read, write";
    permission java.io.FilePermission "${oo.home}/var/logs",
        "read, write";
};
```

Note: You may need to add more permissions depending on the content you run.

To allow the flow to access to resources in the local file system of Central/RAS, you should specify this in the java.policy. For example:

```
grant codebase "file:${oo.home}/var/cache/-" {
    permission java.io.FilePermission
"C:\\users\\cathy\\foo.bat", "read, write, execute, delete";
    permission java.io.FilePermission "C:\\users\\cathy\\-",
"read,write,execute,delete"; // Recursive Example
    permission java.io.FilePermission "C:\\users\\cathy\\*",
"read,write,execute,delete"; // Flat Example
        .....
};
```

Adding Java Security Manager

To avoid a possible scenario where user creates a flow which shuts down the localhost server, apply the specifications from ["Preventing Flows from Accessing the Central/RAS Local File System" on the previous page](#)

The following sections describe how to set up the Java policy file for standalone RAS and Central embedded worker.

Configuring RAS to Add Java Security Manager

To activate security manager on a standalone RAS:

1. In the `%RAS_HOME%\conf` folder, create a text file named `ras.policy`, add the following content:

```
grant codeBase "file:${oo.home}/lib/score-worker-manager-impl.jar"  
{  
  permission java.lang.RuntimePermission "exitVM.75";  
}
```

2. In the `ras-wrapper.conf` file located in the `%RAS_HOME%\conf` folder, add the following wrapper:

```
wrapper.java.additional.<next_index>=-Djava.security.policy="%RAS_  
HOME%/conf/ras.policy"
```

3. Restart the RAS.

Configuring Central Embedded Worker to Add Java Security Manager

To activate security manager on a Central embedded worker:

1. Add the following content before **System Code Permissions** in the `%CENTRAL_HOME%/tomcat/conf/catalina.policy` file.

```
grantcodeBase "file:${oo.home}/tomcat/webapps/oo/WEB-INF/lib/score-  
worker-manager-impl.jar"  
{  
  permission java.lang.RuntimePermission "exitVM.75";  
}
```

