Hewlett Packard
Enterprise

Operations Orchestration

Software Version: 10.70
Windows and Linux Operating Systems

Central User Guide

Document Release Date: November 2016
Software Release Date: November 2016

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

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.

About this PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format. Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online help.

Contents

Welcome to HPE OO Central User Guide

HPE OO Central is a web-based application used for promoting flows, running and scheduling flows, administering the system, and extracting and analyzing data resulting from flow runs.

Major Tasks and Personas

The main personas who work with Central are the Operations Administrator (Op Admin) and the End User. However, some of the initial setup tasks are performed by the System Administrator (Sys Admin).

Operational Administrator (Op Admin)

	The Op Admin is responsible for the day to day operation of HPE OO. This includes configuration, maintenance, content pack promotion, and setting permissions for HPE OO users. The Op Admin troubleshoots any failures in the production environment and either solves the problem or routes it to the System Administrator, to HPE Support, or to the Flow Author (depending on the issue).

The Op Admin uses Central to:

- Promote content packs - see "Promoting a Content Pack – Overview" on page 123
- Run and monitor flows - see "Running and Monitoring a Flow – Overview" on page 179

End User

	The End User triggers and monitors flows. The End User can access entitled HPE OO flows directly through Central or indirectly through an embedded web UI in another application

The End User uses Central to:

- Run and monitor flows - see "Running and Monitoring a Flow – Overview" on page 179

System Administrator (Sys Admin)

	The System Administrator is responsible for the HPE OO hardware and software. The Sys Admin installs and patches HPE OO (Central and RASes), and is responsible for the correct functioning of the application from the system perspective, dealing with things such as CPU, memory, and OS environment.

The Sys Admin uses Central to set up the Central configurations:

- Configure the topology – see:

  "Setting Up Topology – Workers and RASes" on page 81

  "Setting Up Topology – Automatic RAS Upgrade" on page 93

- Configure security and authentication – see:

  "Setting Up Security – Roles" on page 32

  "Setting Up Security – Internal Users" on page 65

  "Setting Up Security – LDAP Authentication" on page 42

  "Setting Up Security – LWSSO" on page 77

  "Setting up Security Settings" on page 25

Sample Scenario

	System Admin		Op Admin		Flow Author

Setup

	After the installation of Central, the **system administrator**, configures the LDAP authentication (see "Setting Up Security – LDAP Authentication" on page 42).
	The system administrator sets up the Central roles and maps these roles to the LDAP groups (see "Setting Up Security – Roles" on page 32).
	The system administrator assigns workers to the worker groups. He creates a new group, ”Windows”, for the external worker (see "Setting Up Topology – Workers and RASes" on page 81).
	The operations administrator deploys the base content in the Central servers (see "Deploying and Managing Content Packs" on page 127).

Promotion

	The flow author creates a new content pack in Studio and deploys it to the file system.
	The operations administrator receives an email from the flow author, saying that a new content pack is ready for deployment.
	The operations administrator deploys the new content pack to the Central Staging server (see "Deploying and Managing Content Packs" on page 127).
	The operations administrator maps the Central user accounts to the system accounts in the content pack (see "Setting Up Configuration Items for a Content Pack" on page 164).
	The operations administrator sets up the system properties in the content pack (see "Setting Up Configuration Items for a Content Pack" on page 164).
	The operations administrator verifies that the Windows alias was automatically assigned to the worker group and maps the other group aliases to actual worker groups (see "Setting Up Configuration Items for a Content Pack" on page 164).

	The operations administrator takes a look at the flow library and verifies that all the flows are there and that when she selects a flow, she can see the flow meta data (see "Managing the Flow Library" on page 149).
	The operations administrator sets the permissions for the flows in the content pack (see "Managing the Flow Library" on page 149).

Running and monitoring flows

	The operations administrator triggers one of the flows in the content pack (see "Running a Flow" on page 181).
	The operations administrator navigates to the **Runs** view to see the flow status (see "Tracking and Managing Flow Runs" on page 215).
	The operations administrator pauses, resumes, or cancels the flows, if required (see "Tracking and Managing Flow Runs" on page 215).
	The operations administrator drills down to the flow details and tests the flow, to see if there are any problems (see "Testing and Troubleshooting a Flow Run" on page 229).

Updating the content pack

	The operations administrator identifies a problem with the content pack. It is too large, and she wants it to be split into two smaller content packs.
	The operations administrator sends an email to the flow author, asking him to split the content pack into two.
	The flow author splits the content pack into two smaller content packs and sends them to the operations administrator.

Final deployment

	The operations administrator deploys the new content packs to the Staging server and verifies that there are no other problems.
	The operations administrator deploys the content packs to the Production server.

Logging into HPE OO

If authentication has been set up for HPE OO, the Logon screen is shown when HPE OO is started.

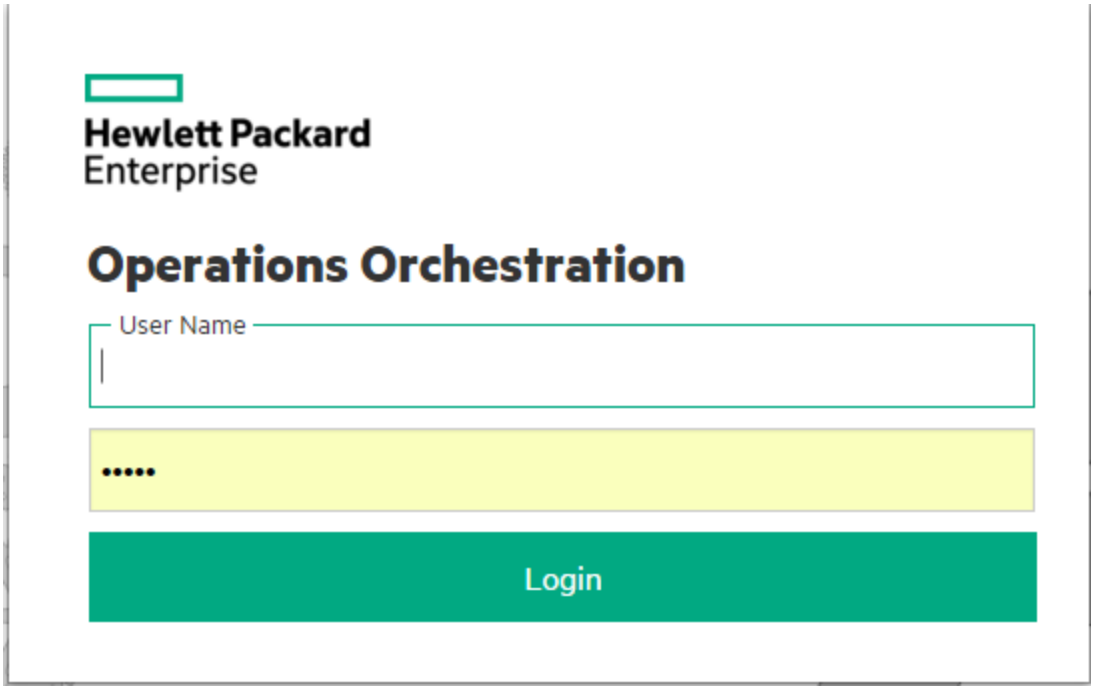

1. Enter your user name and password.

   > **Note:** You cannot log in with a user name that includes a colon.

2. If multiple domains are available, you also need to select a domain. The domain that you select will remain the default choice for the browser that you're using.

   > **Note:** Your current domain is selected by default.

   There is no choice of domain if:

   - Only internal users have been set up
   - There is only one domain and this has been set as the default

     > **Note:** If there is a single domain, but it has not been set as the default, you will have a choice between this domain and the internal users.

> If there are multiple domains, you must select a domain name, unless you are logging in as an internal user or if you belong to the default LDAP.

3. If the option to capture logged-in user credentials has been enabled, the Login screen includes a warning message, saying that HPE OO will capture your credentials and may transfer them over the network.

   For information about how to set up the capture of logged-in user credentials, see "Setting up Security Settings" on page 25.

4. Click **Login**.

If authentication has not been enabled, the Logon screen is not shown. For information about how to set up authentication for HPE OO, see "Setting up the System Configuration" on page 24.

Navigating HPE OO Central

HPE OO Central is divided into four workspaces, so you can perform end-to-end tasks from the same place. A workspace is a unit that holds all the screens that belong to the same end-to-end task group.

- **Dashboard** – used to display the system's ROI, and analyzed flow aggregation. For users with an Op Admin  role.
- **Run Management** – used for running flows, monitoring runs, scheduling runs, and troubleshooting runs. Can be used by users with an End User  or Op Admin  role.
- **Content Management** – used for promotion tasks, such as deploying new content, setting permissions on flows, setting up configuration items, and rolling back to earlier versions of content packs. For users with an Op Admin  role.
- **System Configuration** – used to configure topology, and set up users, roles, LDAP authentication, LWSSO, security banners, monitor and control the size of the database, and personalize the Central UI appearance. For users with a System Administrator  role.

Dashboard

Click the **Dashboard** button to display the Dashboard workspace.

This workspace provides statistical information about the system (popular flows, result distribution, execution time, and so on) and financial information about ROI (return on investment).

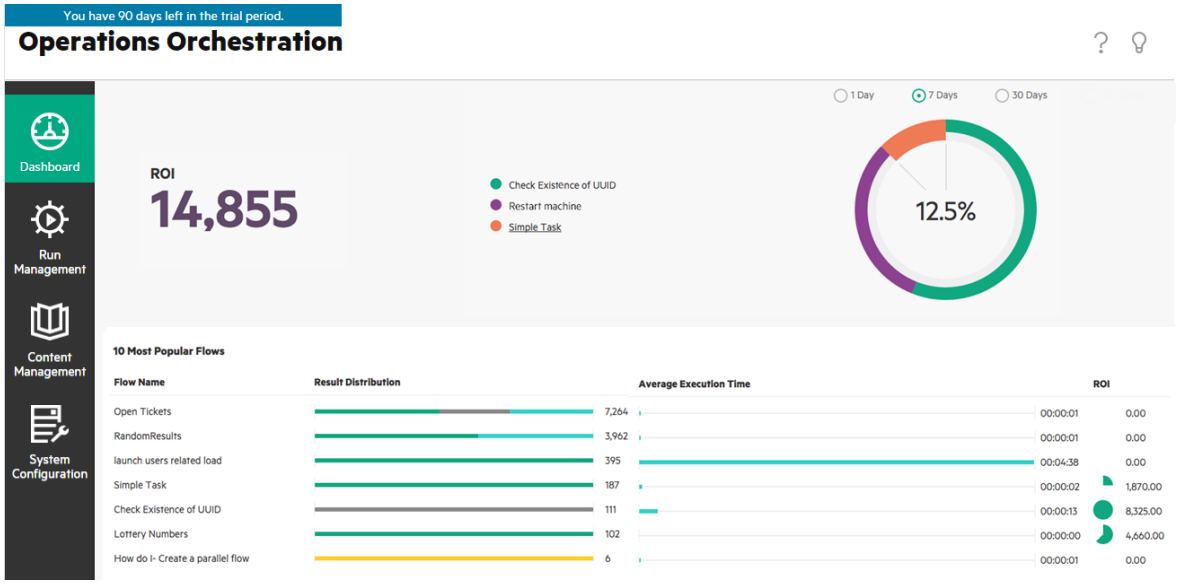

Run Management

Click the **Run Management** button to display the Run Management workspace.

This workspace includes the following modules:

- **Run Explorer** – In the Run Explorer module, the Op Admin or End User can monitor their running flows and the flows that have finished running. You can track flow runs, monitor their progress, and perform actions on flow runs, such as pausing, resuming, and canceling them. If a flow fails and you want to troubleshoot, you can drill down into the run to display detailed information. If you have many flows running at the same time, you can use filters to locate the flow that you need.

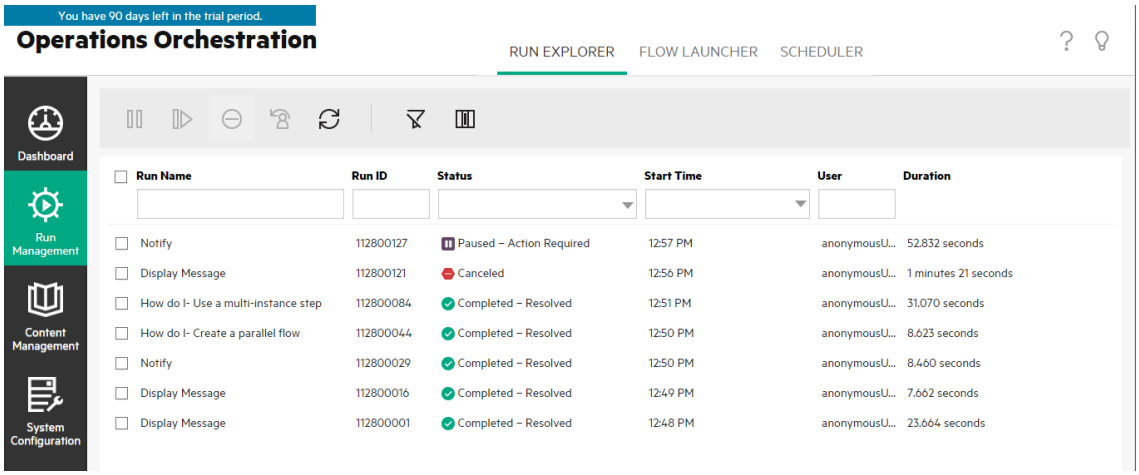

- **Flow Launcher** – In the Flow Launcher module, the Op Admin or End User can browse for a flow, view the flow information, name the run, enter inputs, and run the flow.

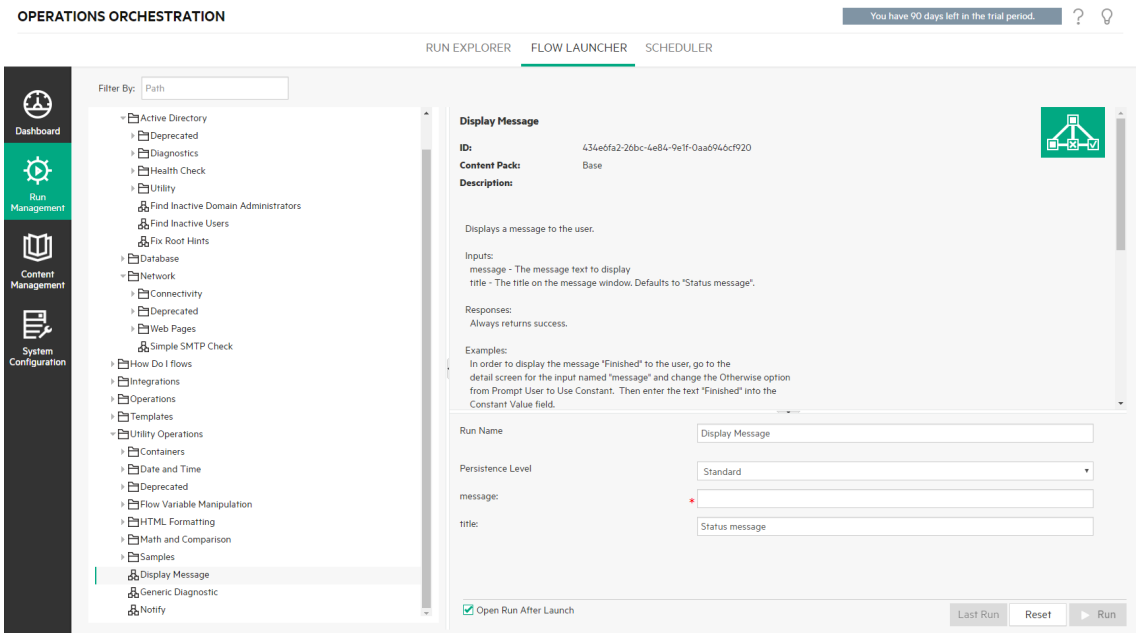

- **Scheduler** – In the Scheduler module, the Op Admin can view, create, and modify flow run schedules.

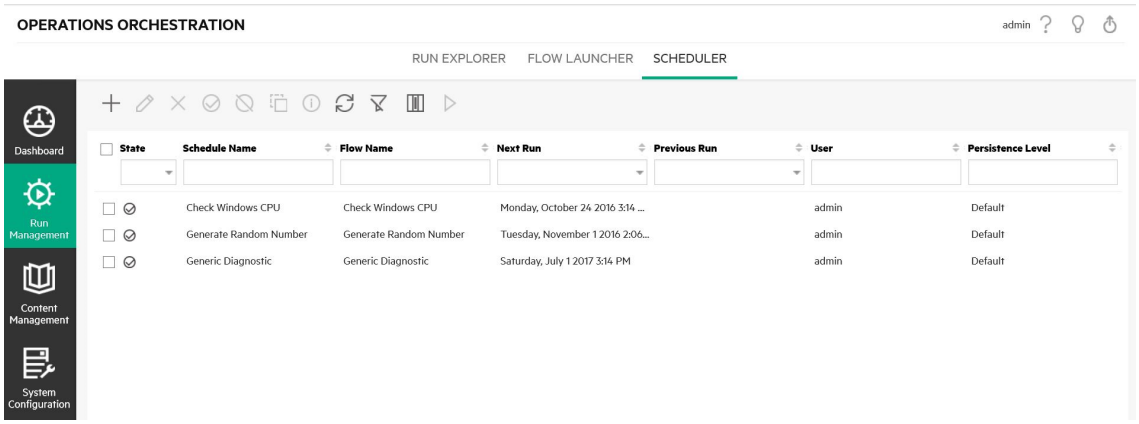

Content Management

Click the **Content Management** button to display the Content Management workspace.

This workspace includes the following modules:

- **Flow Library** – The Flow Library displays the flows that have been deployed. From here, the Op Admin can run or schedule a flow, view a flow's meta data, set the content permissions, or run a flow for testing.

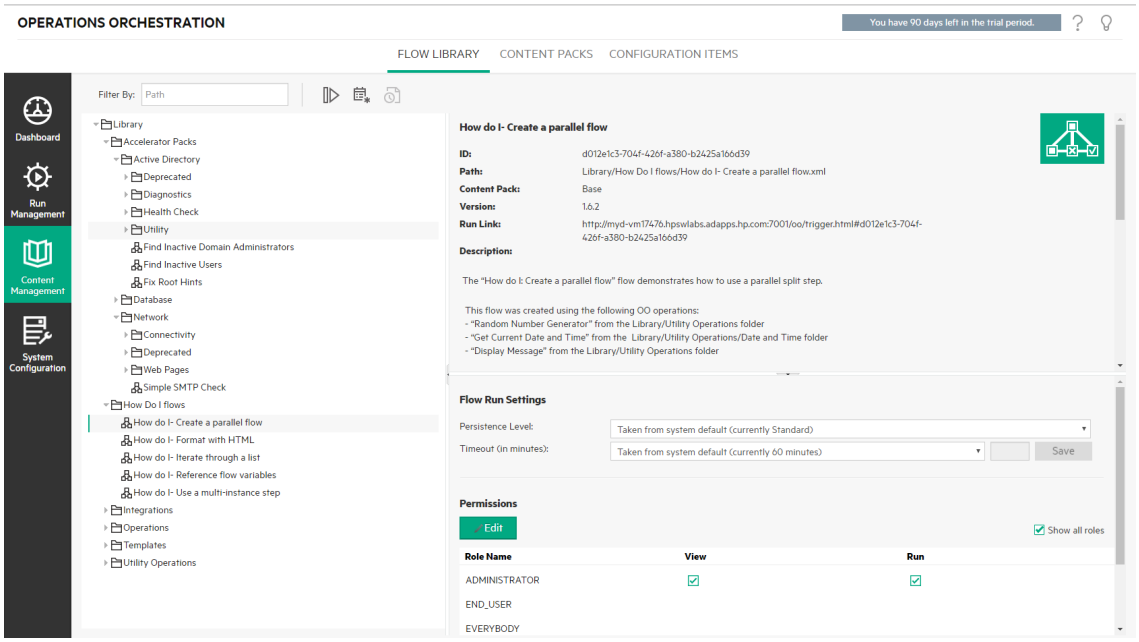

- **Content Packs** - In this module, the Op Admin can view the deployed content packs, filter them, delete them, deploy or roll back content packs, view the details of content packs, compare statistics about deployed content packs, and view changes that were made to them.

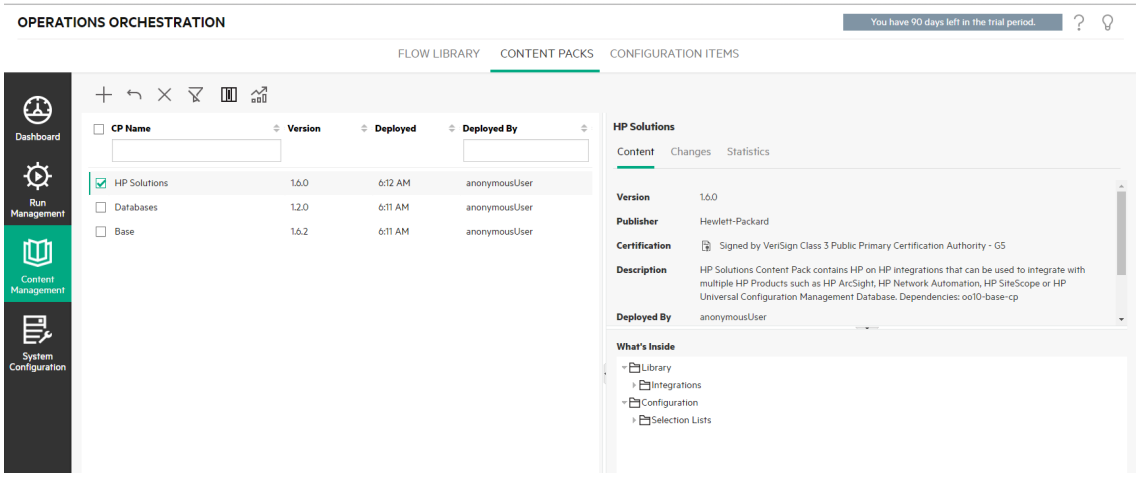

- **Configuration Items** – In this module, the Op Admin can configure the contents of a content pack, in order to adjust it to the environment: mapping group aliases to actual worker groups, mapping system accounts in the content pack to Central user accounts, and mapping system properties.

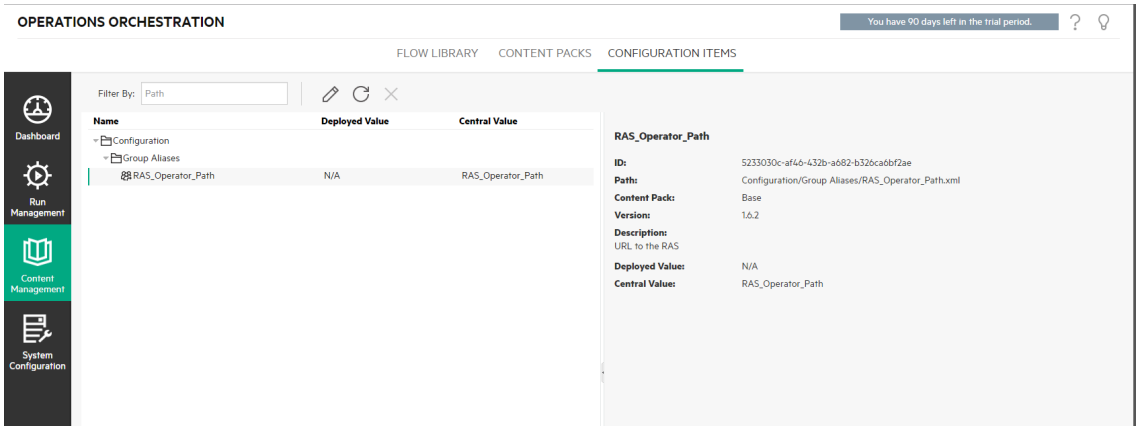
OPERATIONS ORCHESTRATION
You have 90 days left in the trial period.
FLOW LIBRARY
CONTENT PACKS
CONFIGURATION ITEMS
Dashboard
Run Management
Content Management
System Configuration
Filter By:
Path
Name
Deployed Value
Central Value
Configuration
Group Aliases
RAS_Operator_Path
N/A
RAS_Operator_Path
RAS_Operator_Path
ID:
5233030c-af46-432b-a682-b326ca6bf2ae
Path:
Configuration/Group Aliases/RAS_Operator_Path.xml
Content Pack:
Base
Version:
1.6.2
Description:
URL to the RAS
Deployed Value:
N/A
Central Value:
RAS_Operator_Path

System Configuration

Click the **System Configuration** button to display the System Configuration workspace.

This workspace contains the following modules:

- **Security** – Enables the System Administrator to set up roles for HPE OO and to apply these roles to internal users or to users in an authentication system such as LDAP or LWSSO. This is also where the System Administrator can enable authentication and auditing, and configure a security banner.

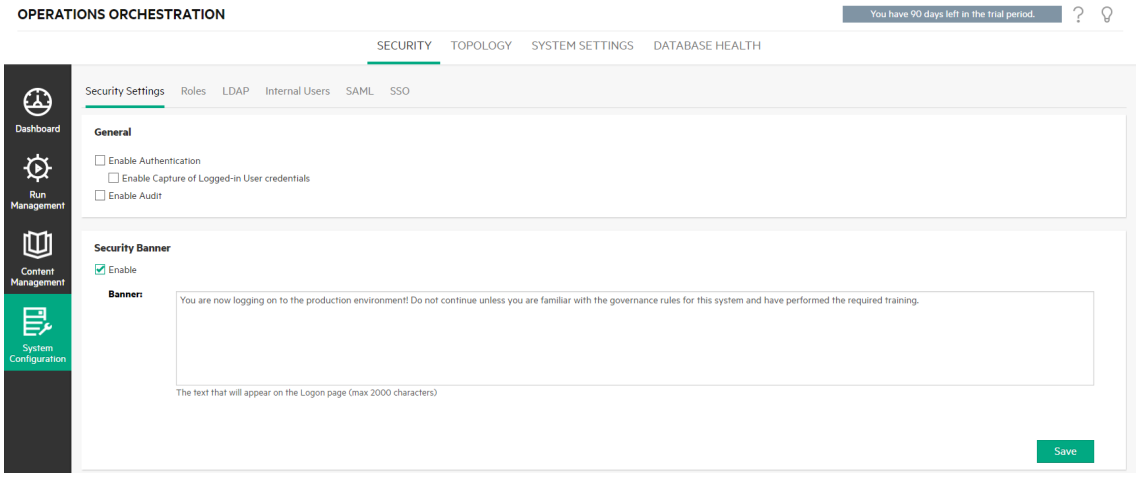

- **Topology** – Enables the System Administrator to configure workers, worker groups, and RASes, and to set up an automatic upgrade process for RASes.

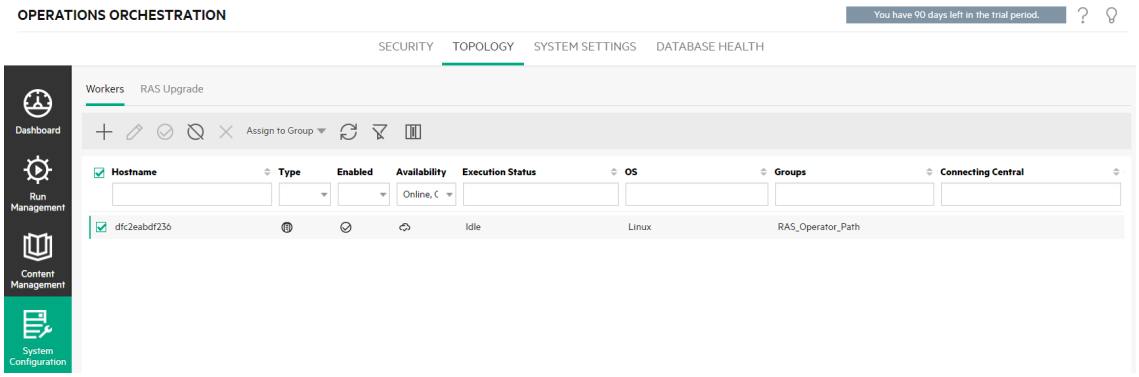

- **System Settings** – Enables the System Administrator to set defaults and customizations, and configure a forward proxy, reverse proxy, or load balancer. It also enables the System Administrator

to install an HPE OO license.

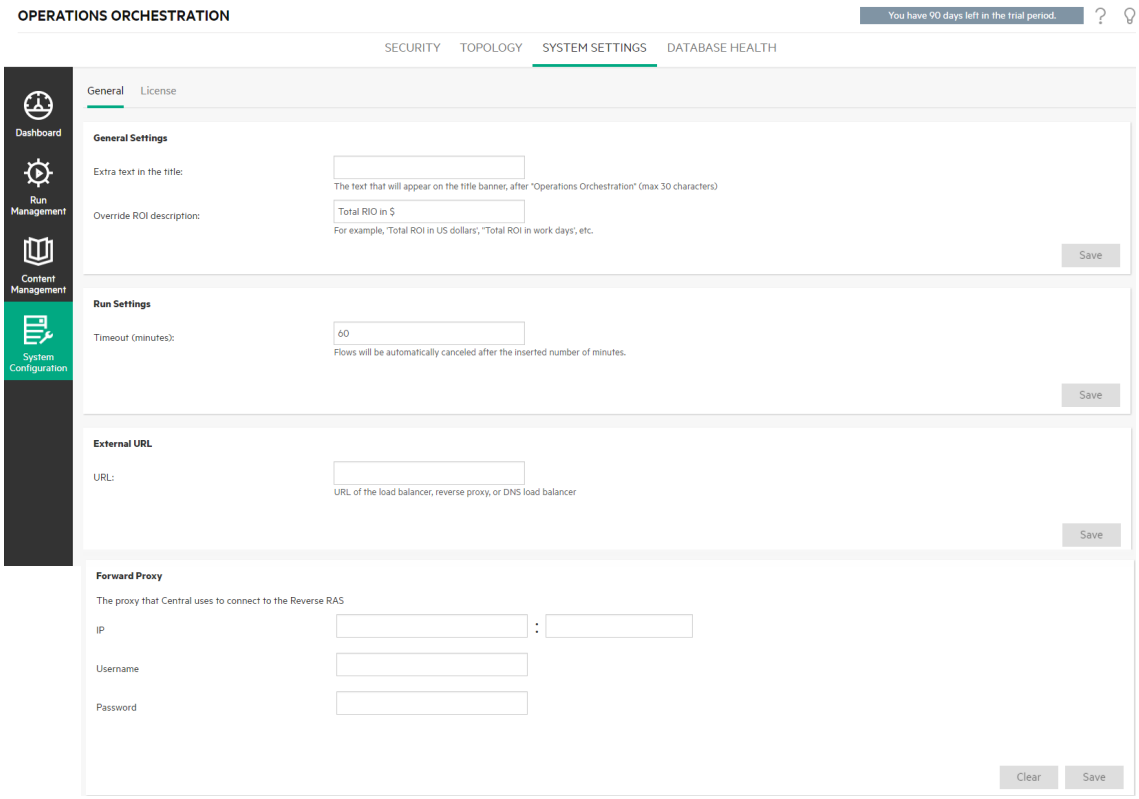

- **Database Health** – Enables the System Administrator to monitor the size of the database and control how much information is saved to the Run Log.

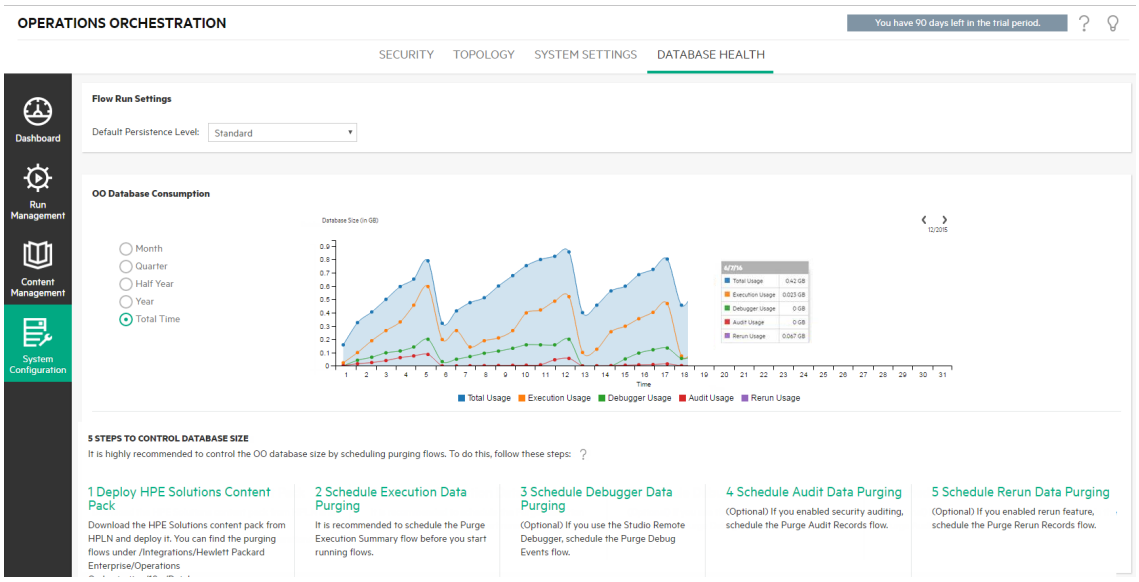

Viewing the Workspaces

It is possible to set up the HPE OO roles so that users who have been assigned particular roles will only be able to access the workspaces that are relevant to their role.

For example, you can set up roles so that:

- Users who belong to the **End User** role will see only the Run Management workspace.
- Users who belong to the **Promoter** role will see the Content Management workspace and the Run Management workspace.
- Users who belong to the **System Administrator** role will see the System Configuration workspace and the Run Management workspace.

For more information about roles, see "Setting Up Security – Roles" on page 32.

Adjusting the Display of Panes in the Workspace

Use the sliders between panes, to adjust the display:

- Drag the edge of a slider to adjust the height or width of a pane.
- Use the open/close button in a slider to toggle between hiding and displaying a pane.

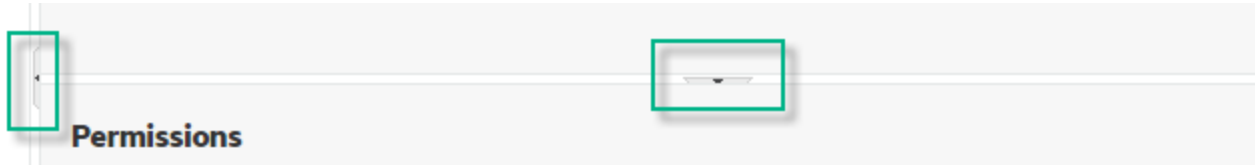

Working with the Dashboard

The Dashboard provides statistical information about the system (popular flows, result distribution, execution time, and so on) and financial information about return on investment (ROI).

The Dashboard shows information about all flows that the logged-in user is entitled to see.

The Dashboard enables the Op Admin to analyze the flow statistics, and compare the performance of the top ten flows.

> **Note:** The Dashboard does not include flows that were run using the Remote Debugger in Studio. This ensures that the statistics and ROI calculations in the Dashboard will not be affected by test runs.

The time range of the data in the Dashboard appears in the top right corner. This time range can be set to the last day, the last seven days, or the last thirty days.

> **Note:** It is also possible to perform a wider range of Dashboard tasks using APIs:
>
> - Set the date range with a wider range of possible date ranges.
> - Display a larger number of flows
> - Sort and filter data
>
> For more information, see "Dashboard" in the *HPE OO API Guide*.

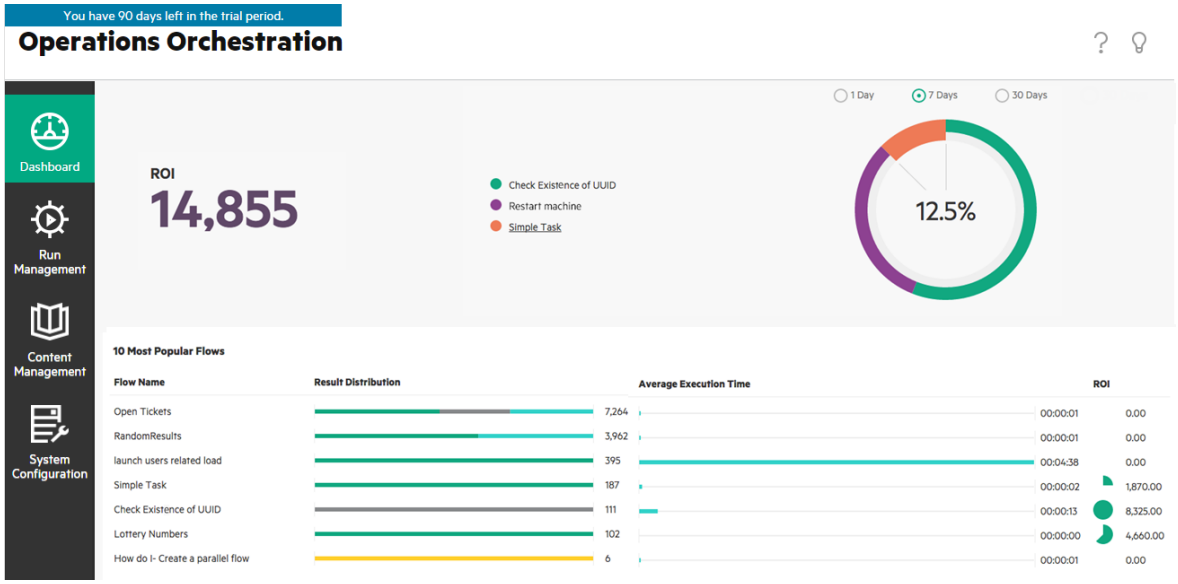

For information about how to interpret the Dashboard, see "Reference Material" on the next page.

> **Note:** The **Dashboard** tab is only visible if you have been assigned a role with the **View Dashboard** permission. For information about setting up roles, see "Setting Up Security – Roles" on page 32.

If you do not have ROI values assigned in your flows, the total ROI value will be 0.

> **Note:** You can personalize the Dashboard by adding text to the ROI title, in order to help users understand the meaning of the ROI values. For example, you can tell them if ROI is calculated in dollars or in work days. For more information, see "Setting Up the System Settings - General Settings" on page 100.

To see a movie on how to set an ROI value for your content, click the **Play** button under the movie thumbnail in the Dashboard.

What do you want to do?

Customize the Dashboard time range

In the dashboard, you can display information about all the flow runs that you are entitled to see, over the past day (the last 24 hours starting from now), the past week (the last 7 days starting from today), or the past 30 days (starting from today).

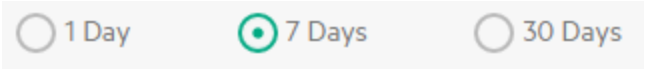

If you do not adjust the date range, the default is seven days.

> **Note:** It is also possible to set the date range using APIs, with a wider range of possible date ranges. For more information, see the *Operations Orchestration API Guide*.

1. Click the **Dashboard** button to display the Dashboard workspace.
2. In the top right corner, select the option with the number of days that you want to include in the date range:
   - **1** - to display information about runs from the past day
   - **7** - to display information about runs from the past 7 days
   - **30** - to display information about runs from the past 30 days

Extract the ROI Information by API

You can extract the relevant ROI information by API, as a basis for periodical reports for the business manager.

For more information, see "Get Statistics" in the *HPE OO API Guide*.

.

Reference Material

System ROI

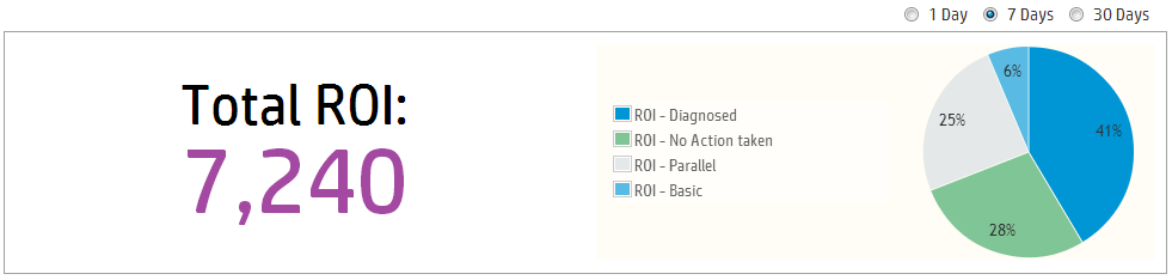

GUI item	Description
Time range	All the data in the Dashboard is relevant to the specified time range in the top right corner. To adjust the date range, select one of the options : • **1 Day** - to display information about runs from the past day • **7 Days** - to display information about runs from the past 7 days • **30 Days** - to display information about runs from the past 30 days
Total ROI	This value is the sum of the ROI values of all performed transitions in the defined time range, which the logged on user is entitled to see.
Main Contributors pie chart	Displays the flows that were the major contributors to the ROI, with their relative contributions expressed in a pie chart.

10 Most Popular Flows

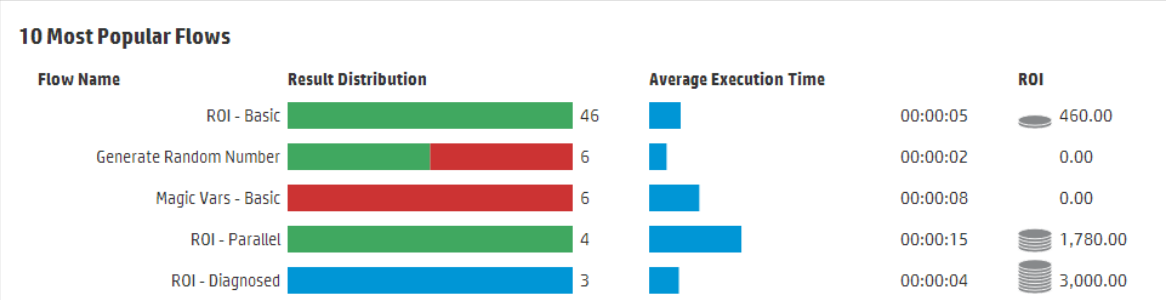

GUI item	Description
Results Distribution bar graph	For each of the top ten flows, displays the number of finished runs. Finished runs means those with a status of success, error, failed to complete, diagnosed, or no action taken. This does not include flows that are paused, running, canceled, pending cancel, or pending pause. The bar graph displays the relative distribution of runs according to status. The different statuses are expressed by the following colors: • **Green**: Completed - Resolved • **Red**: Completed - Error • **Gray**: Completed – No Action Taken • **Blue**: Completed - Diagnosed • **Yellow**: Failed to Complete The colors are the same as the colors of the status icons in the Run Explorer. **Tip** To identify the status indicated by a color, roll the cursor over a colored bar to display a tool-tip.
Average Execution Time bar	Displays the average duration of each flow, based on the finished runs for that flow over the time that was set as the date range. The bar displays the duration time of the flow, relative to the other flows in the top ten.
ROI column	Displays the ROI value of each flow. The coins represent the relative ROI score of the flow compared to the other flows in the top ten.

Setting up the System Configuration

Setting up the system configuration is usually done by the system administrator, after HPE OO is installed.

Setting up Security Settings

The administrator can enable or disable the following security settings for Central:

- Authentication
- Capturing of logged-in user credentials
- Auditing
- A customized security banner

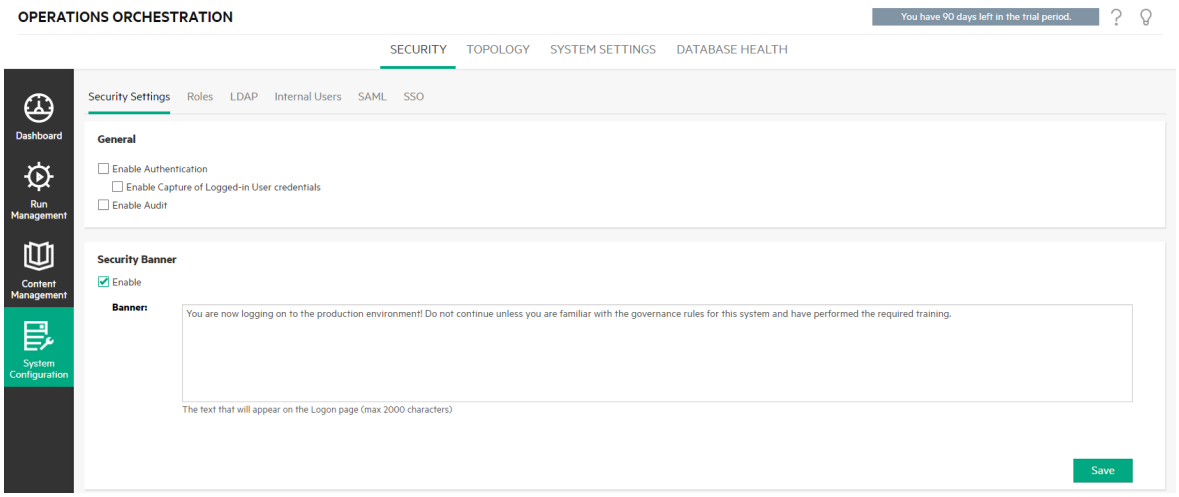

> **Note:** The **Security > Security Settings** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the settings if you have been assigned a role with the **Manage Security Configuration** permission.

Enabling Authentication

The **Enable Authentication** check box is available under the **Security > Security Settings** tab. You can use this to enable or disable the authentication functionality in Central.

> **Note:** You must enable authentication immediately after installation. If authentication is not enabled, all users will be able to access all tasks.

Only users with administrator credentials can enable authentication in Central. This is to prevent a user from becoming locked out of the system.

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.

Enabling Auditing

HPE OO gives you 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.

The **Enable Audit** check box is available under the **Security > Security Settings** tab. You can use this to enable or disable the auditing functionality in Central.

Once auditing has been enabled, a user who has been granted **View Audit** permission can retrieve an audit trail.

Notes:

- By default, the **Enable Audit** check box is not selected.
- Audit events are persisted in the database.
- Currently, it is only possible to retrieve an audit trail via APIs.

  For more information and for the details of the audit events, see the HPE OO *API Guide*.

Enabling the Capture of Logged-in User Credentials

You can enable HPE OO to capture the credentials of the logged-in user. These credentials will then be used in flows that have inputs assigned from logged-in user credentials.

- When this option is **not enabled** (this is the default), the flow will prompt for a password at the relevant step.
- When this option is **enabled**:
  - If a user starts or resumes a flow that includes logged-in credentials as an input parameter, the credentials of the logged-in user will be captured in the database and in memory (in a protected

manner) and used during the flow run. This enables you to make the flow run automatically.

> **Note:** The logged-in credentials that are captured are available to Central only in the context of this session. When the user logs out, the credentials are not saved in the database.
>
> The logged-in user name will contain the domain unless the user belongs to the default LDAP or the user is an internal user and a default LDAP is configured in the system.

- When a user logs in to Central, a message appears, warning that the user's credentials might be captured and used during flow runs, if this is required by the runs. When debugging a Central remotely, the Central login page in Studio displays this warning.

When to use this?

This feature is designed for backward compatibility with HPE OO 9.x and is not recommended for new flows. The best practice for new flows is to pass credentials as inputs.

Limitations

There are limitations to using this option:

- Capturing of logged-in user credentials in not supported on scheduled runs
- Capturing of logged-in user credentials is supported only for users logging to HPE OO via the default LDAP or an internal user repository

> **Note:** It is also possible to enable this option by using OOSH to set the value of the **record.loggedin.user.credentials** system property to **true**, as follows:
>
> ```
> ssc --key capture.loggedin.user.credentials --value true
> ```
>
> For more information, see the *HPE OO Shell (OOSH) User Guide*.

Warning! If you enable the capture of the logged-in user credentials, this causes these credentials to pass through the network. Therefore, we highly recommend to use a secured network in order to protect the user credentials.

> **Note:** When this setting is changed, if auditing is enabled, an audit log is created.

Setting Up a Security Banner

You can configure a security banner to appear before the Logon screen. This banner may be useful if you want to warn users about security rules and issues.

When a security banner is enabled, it appears as a pop-up dialog box before the Logon screen, and the user needs to click an **OK** button to continue.

What do you want to do?

Enable authentication

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.
3. Select the **Enable Authentication** check box to display the Enable Authentication dialog box.
4. In the Enable Authentication dialog box, enter your administrator user name and password, and click **Yes**.

> **Note:** If a choice of LDAP domains has been set up, you will also need to select the domain for the authentication.

Disable authentication

1. After the authentication functionality has been enabled, clear the **Enable Authentication** check box.
2. In the confirmation dialog box, click **Yes**.

Enable auditing

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.
3. Select the **Enable Audit** check box.
4. In the confirmation dialog box, click **Yes**.

Disable auditing

1. After the auditing functionality has been enabled, clear the **Enable Audit** check box.
2. In the confirmation dialog box, click **Yes**.

Enable the capture of logged-in user credentials

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.
3. Select the **Enable capture of logged-in user credentials** check box.
4. In the confirmation dialog box, click **Yes**.

Disable the capture of logged-in user credentials

1. If the capture of logged-in user credentials functionality has been enabled, clear the **Enable capture of logged-in user credentials** check box.
2. In the confirmation dialog box, click **Yes**.

Add a security banner

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.
3. In the **Security Banner** section, select the **Enable** check box.
4. In the **Banner** text box, enter the text that you want to appear on the banner. The maximum length of this text is 2000 characters.
5. Click **Save** to save the banner.

Edit a security banner

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.
3. Make the required changes in the **Banner** text box.
4. Click **Save**.

Remove a security banner

If you clear the **Enable** check box, but do not delete the text from the **Banner** text box, you can save this text for later use.

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Click the **Security** tab and then click the **Security Settings** tab beneath it.

3. Clear the **Enable** check box in the **Security Banner** section.
4. Click **Save**.

Reference Material

Security > Enable Authentication

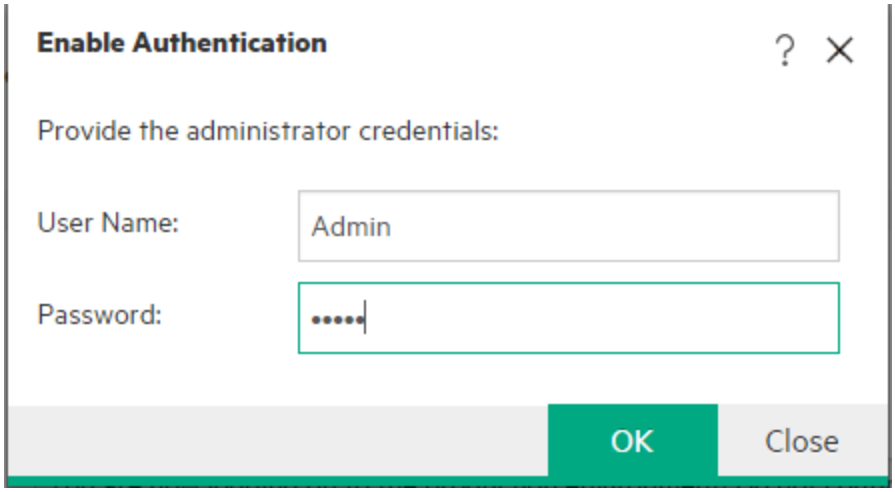

GUI item	Description
User Name	Enter your administrator user name. **Note:** If a choice of LDAP domains has been set up, you will also need to select the domain for the authentication.
Password	Enter your administrator password.
OK	Click to save the credentials and enable the authentication.

Security > Security Settings

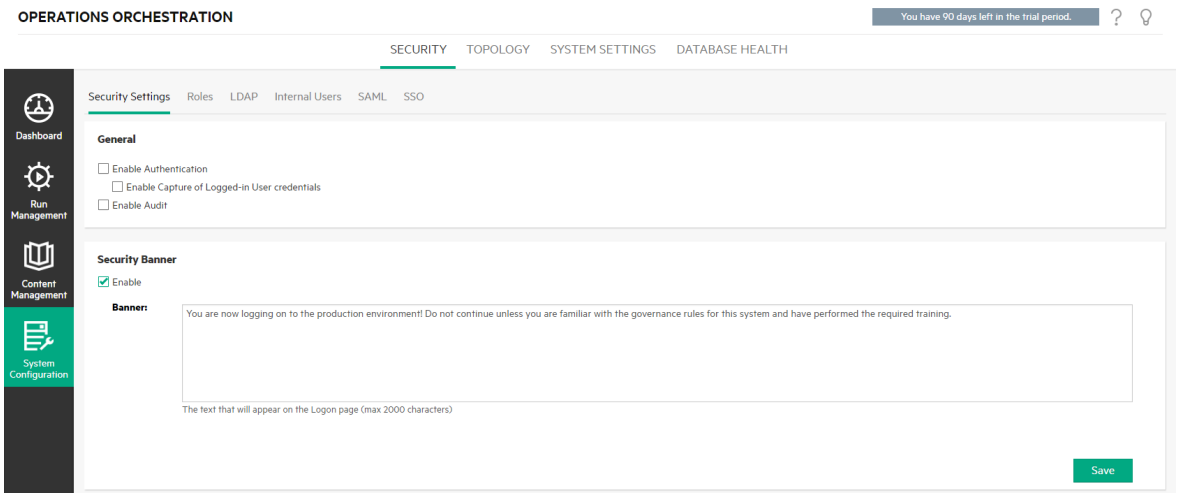

GUI item	Description
General section	
Enable Authentication check box	Select the check box to enable the user authentication functionality in HPE OO. If this check box is cleared, there is no authentication and all users can access all tasks.
Enable Capture of Logged-in User Credentials	Select the check box to enable HPE OO to capture the credentials of the logged-in user.
Enable Audit check box	Select the check box to enable the auditing functionality in HPE OO. Users who have been granted **View Audit** permission will be able to retrieve an audit trail.
Security Banner section	
Enable check box	Select the check box to enable the security banner, so that it is displayed in the Logon screen.
Banner text box	Enter the text for the security banner. The maximum length of this text is 2000 characters.
Save	Click to save the security banner.

Setting Up Security – Roles

 Access to HPE OO Central tasks is managed using Role-Based Access Control (RBAC). RBAC lets the administrator control which actionsHPE OO users are allowed to perform, based on their role.

- A **permission** is a predefined ability to perform a task. Central comes ready with a set of permissions that can be assigned to roles.
- A **role** is a collection of permissions. Roles can be assigned to users.

For example, you could create a role called FLOW MANAGER with permission to run flows, view schedules, schedule flows, manage runs, and create reports. An END USER role could have permission to run flows, but not to schedule them. If a role is set up with no schedule permissions, users with this role do not even see the Scheduler module.

The large number of available permissions enables a refined definition of which parts of the web UI each role sees. The administrator can formulate a dedicated UI experience for each group of users . This enables end users to use Central and only be exposed to the information they need, limiting their ability to harm the system.

> **Note:** The **Role** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the roles if you have a role with the **Manage Security Configuration** permission.

Once the roles are set up, they can be assigned to internal users or mapped to LDAP or SAML users:

- The administrator can map a role to one or more groups, so that LDAP users who are assigned to these groups will be assigned the permissions in the role.
- The administrator can assign a role to an internal user, while setting up the user configuration. For more information, see "Setting Up Security – Internal Users" on page 65.

By default, Central includes the roles ADMINISTRATOR, END_USER, EVERYBODY, PROMOTER, and SYSTEM_ADMIN, corresponding to these roles in Studio.

> **Note:** Role-Based Access Control is not the same as content permissions. If you want to restrict access to data (flows and folders), you will need to configure the permission of that content. This is done as part of the promotion of a content pack. For information about content permissions, see "Managing the Flow Library" on page 149.

What do you want to do?

Create a role

> **Note:** You must have the **Manage Security Configuration** permission in order to create and edit roles.

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **Roles**.
3. Click the **Add** button.

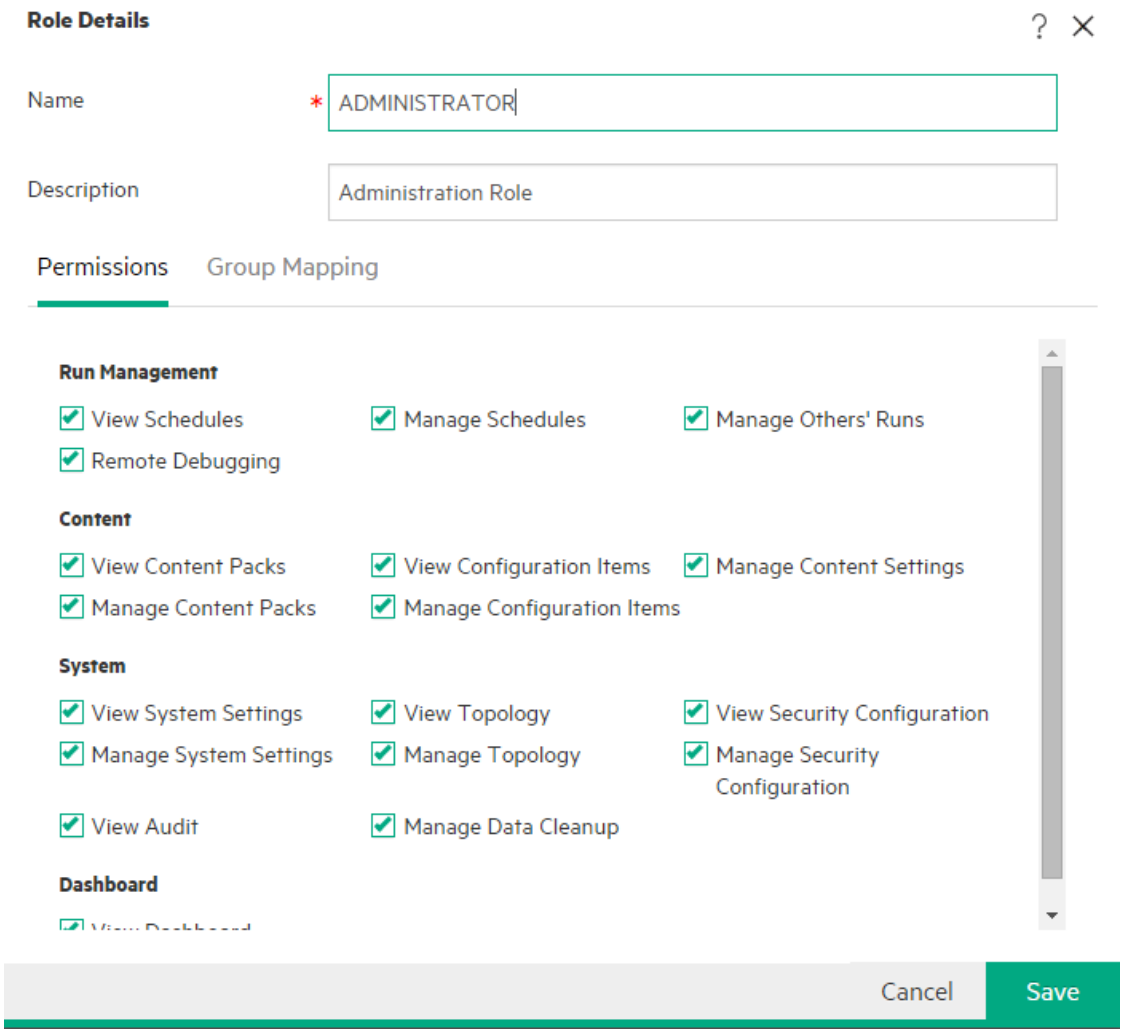

4. In the **Name** box, enter a unique name for the role. The role name cannot have more than 255 characters.
5. In the **Description** box, enter a description of the role. The description cannot have more than 255 characters.
6. Select the check boxes next to the permissions that you want to assign to this role. The permissions are grouped according to the four Central workspaces. If a role is not assigned any permissions in a workspace, users with this role will not be able to enter that workspace.

   Run Management

   - **View Schedules** – Permission to view schedules
   - **Manage Schedules** – Permission to create and modify schedules
   - **Manage Others' Runs** – Permission to modify runs that were triggered by other users
   - **Remote Debugging** – Permission to trigger the Studio Debugger on a remote Central. This permission also gives view and run permissions on all the Flow Library.

   Content

   - **View Content Packs** – Permission to view content packs
   - **Manage Content Packs** – Permission to modify content packs, deploy content, and roll back a deployment.
   - **View Configuration Items** – Permission to view configuration items (for example, group aliases, system accounts, system properties)
   - **Manage Configuration Items** – Permission to modify the values of configuration items (for example, group aliases, system accounts, system properties)
   - **Manage Content Settings** - Permission to modify content permissions (for flows, system accounts, and folders), flow persistence settings, and flow timeout settings.

   System

   - **View System Settings** – Permission to view the Monitoring and System Information reports and to view the system configuration values available through OOSH

     > **Note:** These reports contain information for HP Support:
     >
     > - **/oo/reports/sysinfo** shows the system configuration
     > - **/oo/monitoring** is a performance monitoring dashboard

- **Manage System Settings** – Permission to configure the log level using REST API, to manage the system configuration values available through OOSH, and to manage the Audit configuration
- **View Topology** – Permission to view workers and worker groups
- **Manage Topology** – Permission to enable/disable workers and configure worker groups
- **View Security Configuration** – Permission to view the security configuration. This includes viewing internal users, roles, LDAP authentication, LWSSO, SAML, and Security banner configurations.
- **Manage Security Configuration** – Permission to modify the security configuration. This includes configuring internal users, roles, LDAP authentication, LWSSO, SAML, and the Security banner.
- **View Audit** – Permission to access the audit trail.
- **Manage Data Cleanup** – Permission to use the data cleanup (purging) APIs. For more information, see the *HPE OO API Guide*.

Dashboard

- **View Dashboard** - Permission to view the Dashboard

7. To map this role to an LDAP group, click the **Group Mapping** tab.

   For more information, see Map a role to a group.

8. Click **Save** to save the role.
9. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more information, see "Setting up Security Settings" on page 25.

   > **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.
   >
   > If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

Edit a role

You must have **Manage Security Configuration** permission in order to edit a role.

1. Select **Security** > **Roles**.
2. Select the role that you want to edit and click the **Edit**  button.

3. Make the required changes in the Role Details dialog box.
4. Click **Save**.

Delete a role

1. Select **Security** > **Roles**.
2. Select the role that you want to delete and click the **Delete** button.
3. Click **OK** in the confirmation dialog box.

Map a role to a group

Group mapping is case-sensitive. If you have two groups with similar names (for example, QG_Basic and qg_basic), which you assign to two different roles, a user logged in with one of these roles will not have access to the permissions of the other.

1. Select **Security** > **Roles**.
2. Select the role that you want to map to a group, and click the **Edit** button.
3. In the Role Details dialog box, click the **Group Mapping** tab.

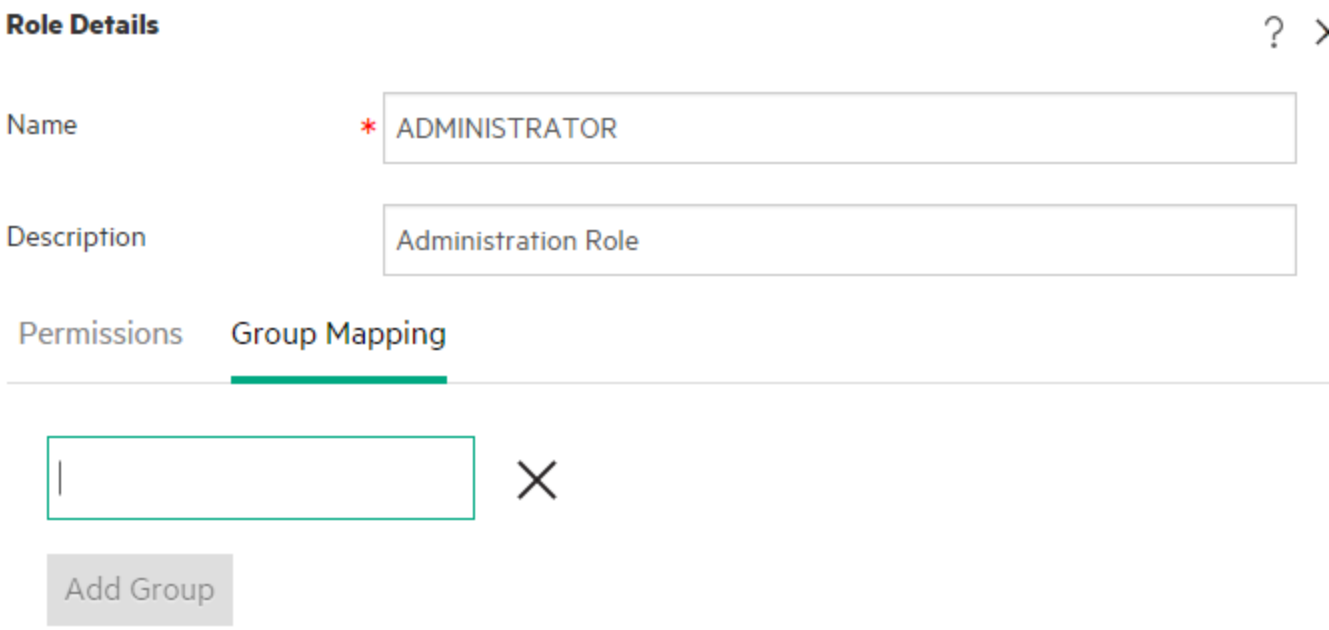

4. Click the **Add Group** button to add a new mapping.
5. In the text box, type the name of the group to which you want to map this role.

6. If required, click the **Add Group** button again to map the role to further groups.
7. Click **Save**.

Delete the mapping of a role to a group

1. Select **Security** > **Roles**.
2. Select the role from which you want to remove the mapping to a group, and click the **Edit**  button.
3. In the Role Details dialog box, click the **Group Mapping** tab.
4. Click the **Delete**  button next to the mapping that you want to delete.

   > **Note:** There is no confirmation dialog box when you delete a mapping. If you change your mind, click **Cancel** to close the Role Details dialog box without saving.

5. Click **Save**.

Assign a role to be the default role

When a role is defined as the default role, it will be assigned to all users that have not been assigned a role. This means that all LDAP users will have the permissions from this default role, in addition to users that are explicitly associated with the role.

> **Note:** If you are setting a default role, it is recommended to set this as the one with the least privileges.
>
> Think carefully about whether you really want to set a default role. When you add permissions to that role, this will affect all LDAP users.

1. Select **Security** > **Roles**.
2. Select the role that you want to specify as the default role.
3. Click the **Set Default Role**  button in the **Roles** toolbar.

> **Note:** You can click the **Set Default Role**  button again to remove the default role functionality from this role.

Reference Material

Security > Roles

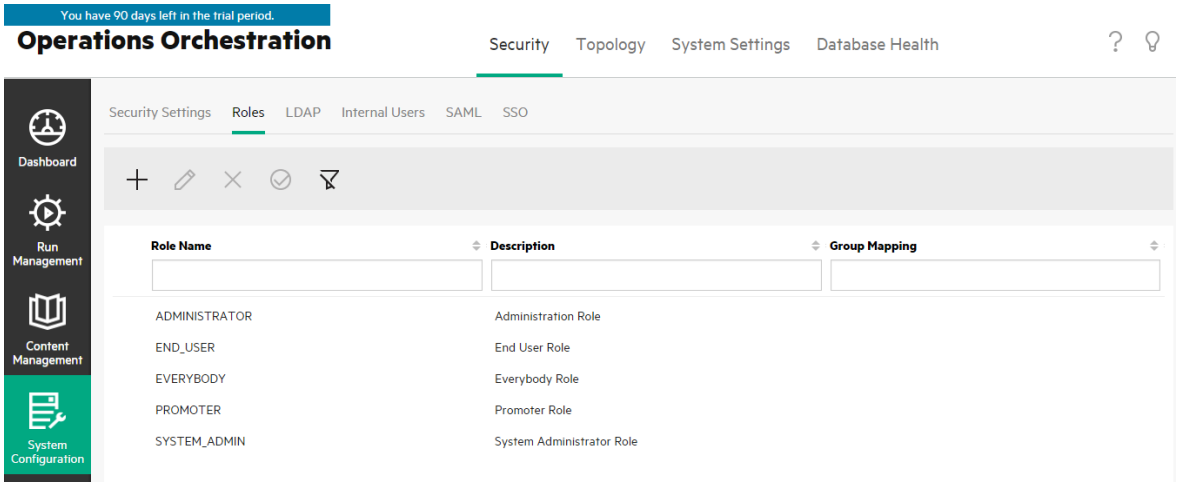

GUI item	Description
Add button	Opens the Role Details dialog box, so that you can add a new role.
Edit button	Opens the User Configuration dialog box, so that you can edit the selected role.
Delete button	Deletes the selected role.
Set Default Role button	Makes the selected role the default role. If the selected role was set as the default role, makes it no longer the default role.
Role Name	Displays the role name.
Description	Displays a description of the role.
Group Mapping	Displays the groups that have been mapped to the role.
Filter Boxes	In the filter boxes at the top of the columns, enter text to filter the displayed roles.

Role Details dialog box > Permissions tab

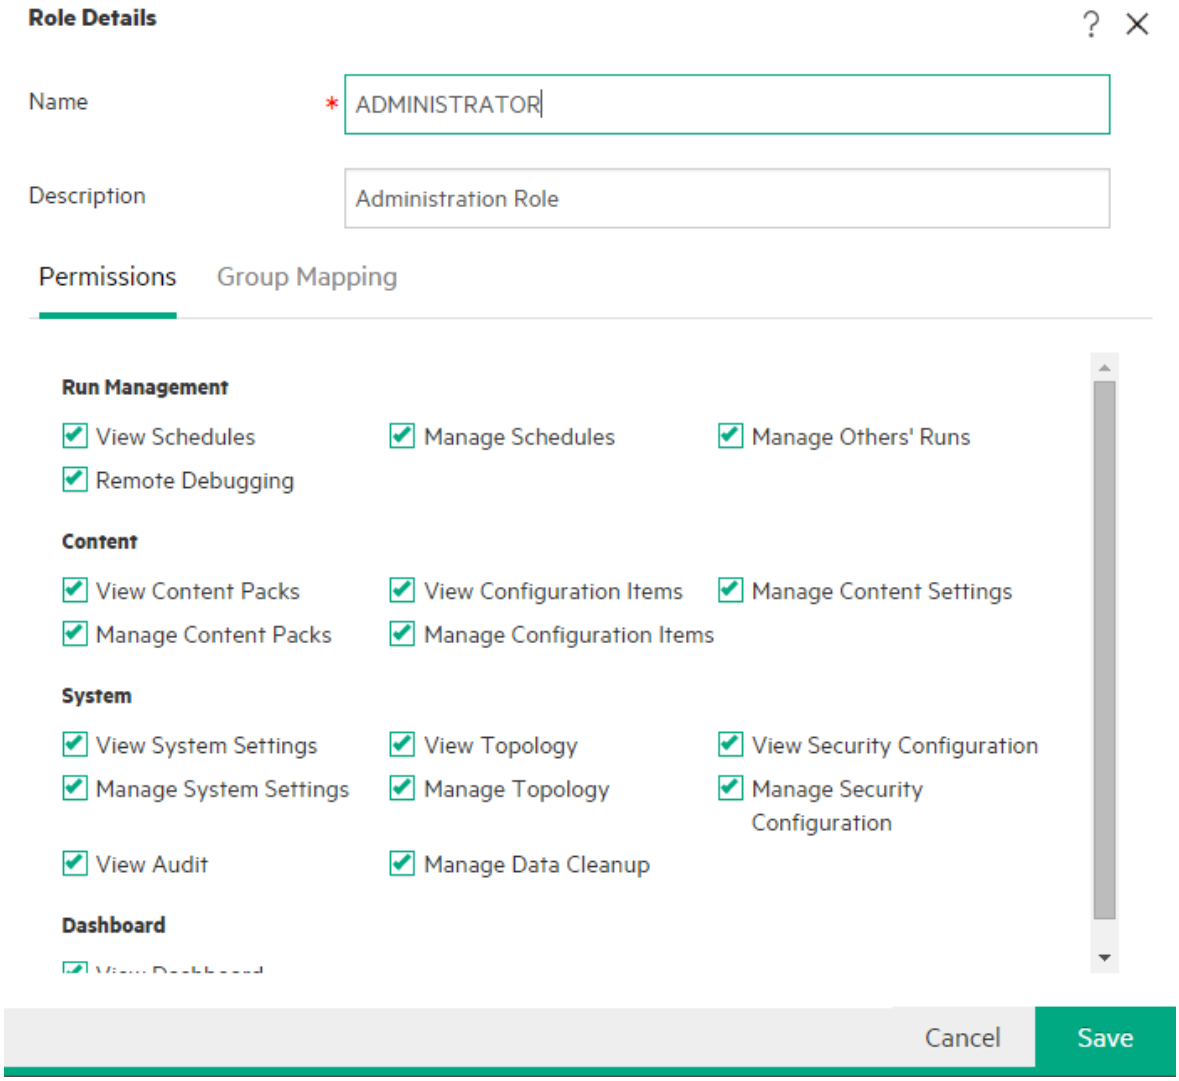

GUI item	Description
Name	Enter a name for the role, of up to 255 characters.
Description	Enter a description of the role, of up to 255 characters.
Run Management	
View Schedules	Permission to view schedules.
Manage Schedules	Permission to create and modify schedules.
Manage Others' Runs	Permission to modify runs that were triggered by other users.
Remote Debugging	Permission to trigger the Studio Debugger on a remote Central. This

	permission also gives view and run permissions on all the Flow Library.
Content	
View Content Packs	Permission to view content packs.
Manage Content Packs	Permission to modify content packs, deploy content, and roll back a deployment.
View Configuration Items	Permission to view configuration items (for example, group aliases, system accounts, system properties).
Manage Configuration Items	Permission to modify the values of configuration items (for example, group aliases, system accounts, system properties).
Manage Content Settings	Permission to modify content permissions (for flows, system accounts, and folders), flow persistence settings, and flow timeout settings.
System	
View System Settings	Permission to view the Monitoring and System Information reports, which contain information for HP Support: • **/oo/reports/sysinfo** shows the system configuration • **/oo/monitoring** is a performance monitoring dashboard Permission to view the system configuration values available through OOSH.
Manage System Settings	Permission to configure the log level using REST API. Permission to manage the system configuration values available through OOSH. Permission to manage the Audit configuration.
View Topology	Permission to view workers and worker groups
Manage Topology	Permission to enable/disable workers and configure worker groups
View Security Configuration	Permission to view the security configuration. This includes viewing internal users, roles, LDAP authentication, LWSSO, SAML, and the Security banner.
Manage Security Configuration	Permission to modify the security configuration. This includes configuring internal users, roles, LDAP authentication, LWSSO, SAML, and the Security banner.
View Audit	Permission to access the audit trail. For more information about auditing, see the *HPE OO API Guide*.
Manage Data Cleanup	Permission to use the data cleanup (purging) APIs. For more information about data cleanup, see the *HPE OO API Guide*.

Dashboard	
View Dashboard	Permission to view the Dashboard.

Role Details dialog box > Group Mapping tab

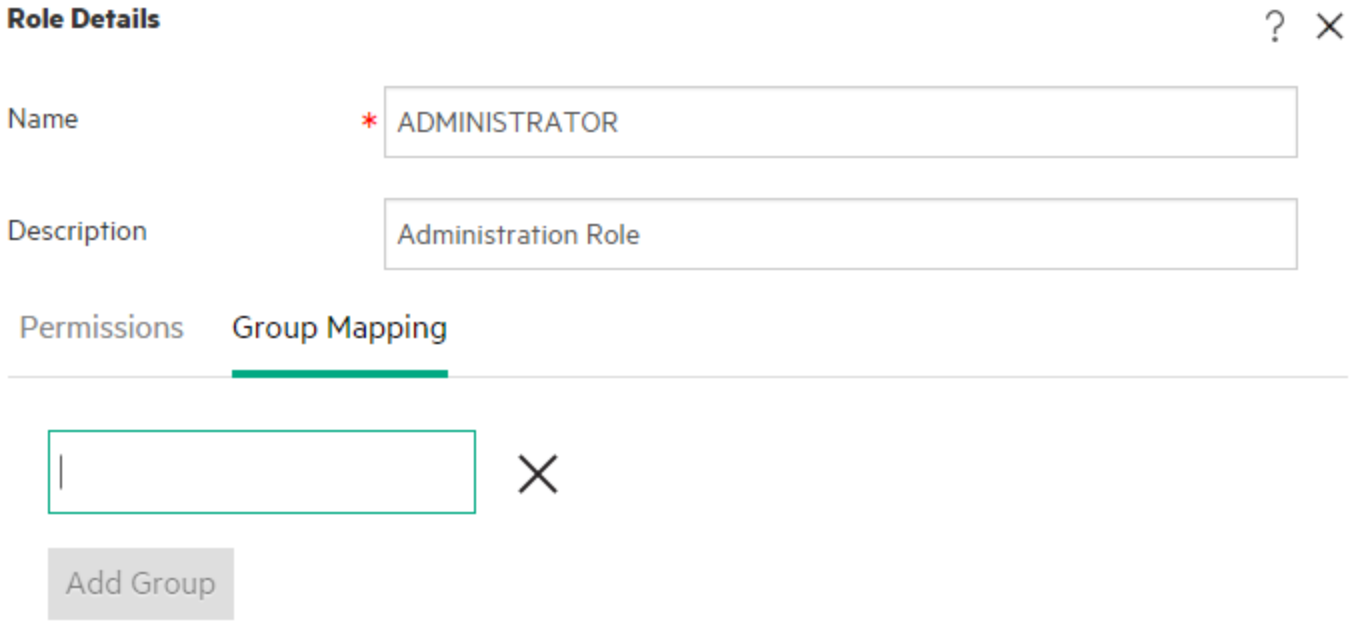

GUI item	Description
Add Group button	Displays a text box so that you can create a mapping from this role to a group.
Text box	In the text box, type the name of the group to which you want to map this role.
Delete ✕ button	Deletes the mapping in the adjacent text box.

Setting Up Security - LDAP Authentication

 If you are using an LDAP authentication service, you can add the LDAP configuration to Central. Users will be able to log in to the system by providing their credentials from their organization. This change is dynamic—there is no need to restart Central for the change to be applied.

The LDAP groups that users belong to can be mapped to HPE OO roles, so that the administrator can control which actions users are allowed to perform in the system. For more information about roles, see "Setting Up Security – Roles" on page 32

If your organization works with multiple LDAP servers, it is possible to configure Central to work with all of them. This includes LDAPs with different schemes and from different vendors. For example, you might have an Active Directory (Microsoft LDAP) implementation for one part of the organization and a Sun One (Oracle LDAP) implementation for another part.

If you set up multiple LDAP configurations on different domains, when users log in to HPE OO, they will need to select from a drop-down list with the active domains. There will be no drop-down list of domains if one of the LDAP configurations is set as the default or if only internal users are configured.

Login Conventions

Users will be able to log in using the following conventions:

- domain\username
- username@domain

The domain name is required unless a user is logging in as an internal user or if they belong to the default LDAP.

It is also possible to log in using the **samAccountName**, **mail**, and **principalName** attributes.

> **Note:** For non-Active Directory configurations, it is also possible to log in with any unique attribute.

Default LDAP

If you flag an LDAP configuration as default, users who belong to this LDAP will be able to log in without having to select a domain.

In earlier versions of HPE OO 10.x, users did not provide a domain when authenticating. The behavior was to go first to the LDAP and after that, to the internal users (if the user was not found). The “default LDAP” feature preserves this behavior for upgraders. When a default LDAP is set, an authentication attempt without a domain will go first to the default LDAP and then to the internal users. See Set the default LDAP authentication configuration.

> **Note:** If you have upgraded from an earlier version of HPE OO 10.x, the LDAP that was configured in the previous version is set as the default.

> **Note:** If you do not have an LDAP service, you can set up internal users, who will log in using their HPE OO user name and password. For more information about this option, see "Setting Up Security – Internal Users" on page 65.
>
> It is also possible to run Central without the user authentication feature. All Central users will be identified as "anonymous" and the system will not require authentication. All Central users will have full administrative access to do anything in the system—triggering, content management, and system configuration.

> **Note:** The **LDAP** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the LDAP authentication if you have a role with the **Manage Security Configuration** permission.

> **Note:** For a comparison between configuring LDAP in HPE OO 10.x and in previous versions, see the HPE OO *Release Notes*.

What do you want to do?

Add an LDAP authentication configuration (not Active Directory)

This procedure is relevant for all LDAP authentication configurations, apart from Active Directory. If you want to create an Active Directory configuration, see Add an Active Directory LDAP authentication configuration.

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **LDAP**.
3. From the **Add** menu, select the type of LDAP configuration that you want to add.

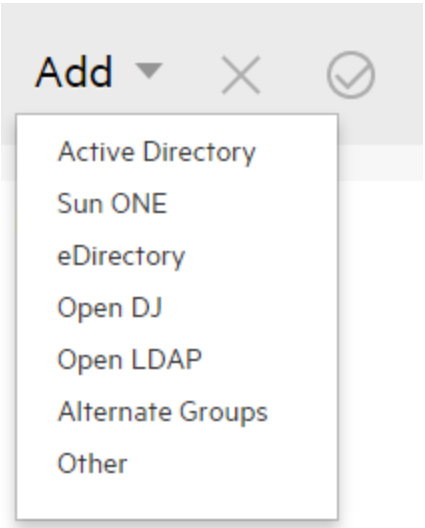

The **LDAP** tab displays the fields relevant to the selected LDAP configuration.

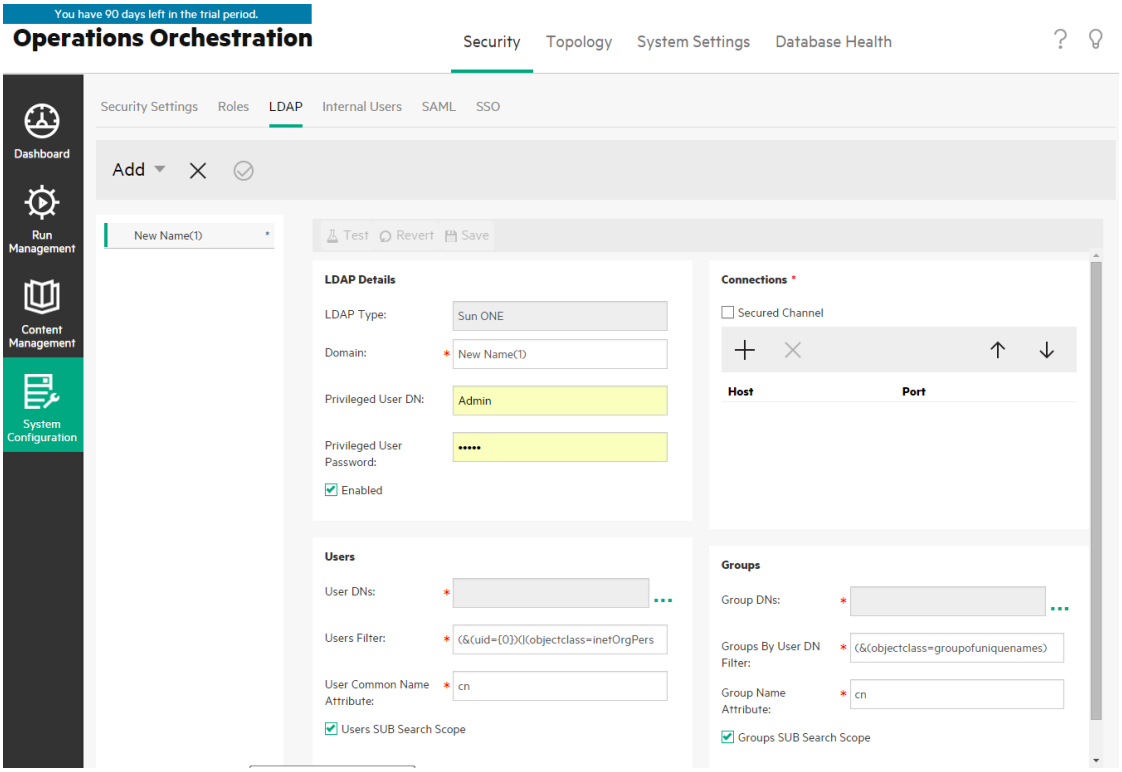

4. In the **Domain** box, enter a name for the domain.

   For example, if you have a single LDAP configuration, you might enter your company name. If you are setting up multiple LDAP configurations, you would give each one a domain that will identify it.

When users log in to HPE OO, they will be able to select from a drop-down list of the active domains. They will be required to select a domain name unless they are logging in as an internal user or if they belong to the default LDAP.

The domain name appears in the LDAP navigation pane to the left.

5. In the **Privileged User DN** box, 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`

   > **Note:** Make sure to use the exact DN for the user with these privileges.

6. In the **Privileged User Password** box, enter the password of the user with search privileges.
7. Select the **Enabled** check box to enable the authentication of this LDAP.
8. In the **Connections** section, select **Secured Channel** to support secured transport layer connection.
9. In the **Hosts** box, enter the IP address or host name of the LDAP server.
10. In the **Ports** box, enter the port number of the LDAP server. This value must be between 0 and 64435.
11. If you have multiple LDAP servers, click the **Add** + button to add a new line, and add the additional host and port.
12. Use the **Up** and **Down**  arrows to move a connection up and down in the list. A connection that is higher in the list has a higher priority in the case of failover.
13. Click the browse button next to the **User DNs** box, to open the User DNs dialog box, in order to define the distinguished names for users.
    - If you know the user DN, click the **Add** + button to add an empty row, and enter the user DN information manually. If you have multiple LDAP hosts, use ";" as the delimiter.

      For example, `ou=people,dc=devlab,dc=ad;ou=people,dc=devlab,dc=com`
    - If you don't know the exact syntax, click **Add Base DN** to fetch a base DN from the LDAP.

      From the list of available root DNs, select the one that is closes to your needs and then refine

it.

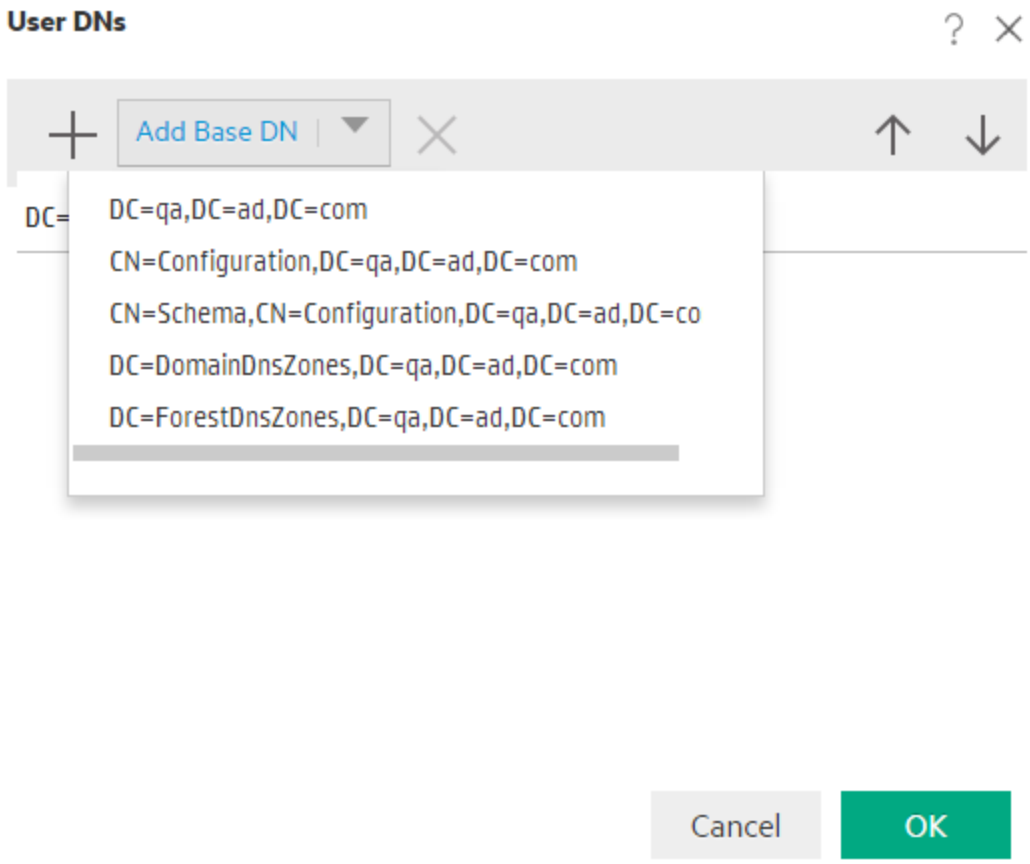

14. Use the **Up** and **Down** arrows  to adjust the order of the user DN rows. The first item in the list is used first.
15. Click **OK** to close the User DNs dialog box and load the information into the **User DNs** box.
16. In the **Users Filter** box, enter the search filter parameters in order to filter users. This filter will be applied on each user DN that was entered in the **Users DNs** box.

    It is recommended to modify the default values, even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute.

    > **Note:** In this filter, `{0}` is a placeholder for the user ID. When a user logs in, `{0}` is automatically replaced with the supplied user name.

17. In the **User Common Name Attribute** box, enter the attribute that defines how the user name will

be displayed on the screen.

> **Note:** In most LDAP implementations, `cn` is the default.

18. Select the **Users SUB Search Scope** check box if you want the users filter to run recursively.

> **Note:** If recursive searching is not required, then it is not recommended to select this check box.

19. Click the browse button next to the **Group DNs** box, to open the Group DNs dialog box, in order to define the distinguished names for groups.

    This step is very similar to defining the distinguished names for users.

20. In the **Groups By User DN Filter** box, enter the search filter parameters in order to filter the groups. This filters the groups of which the given user is a member. The filter is applied on each Group DN.

    It is recommended to modify the default values even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute.

21. In the **Group Name Attribute box**, enter the attribute that defines how the name of the group will be displayed in the UI.

> **Note:** In most LDAP implementations, `cn` is the default.

22. Select the **Groups SUB Search Scope** check box if you want the groups filter to run recursively.

> **Note:** If recursive searching is not required, then it is not recommended to select this check box.

23. Make sure to test the configuration before saving. For more information, see Testing the Configuration.

24. Click **Save** to save the configuration.

> **Note:** Make sure to save your changes before enabling authentication. If you don't save, the authentication will not include your changes!

    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.

25. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more

information, see "Setting up Security Settings" on page 25.

> **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.
>
> If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

Add an Active Directory LDAP authentication configuration

This procedure is relevant for Active Directory configurations only. If you want to create a different kind of LDAP configuration, see Add an LDAP authentication configuration (not Active Directory).

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **LDAP**.
3. From the **Add** menu, select **Active Directory**.

   The **LDAP** tab displays the Active Directory view.

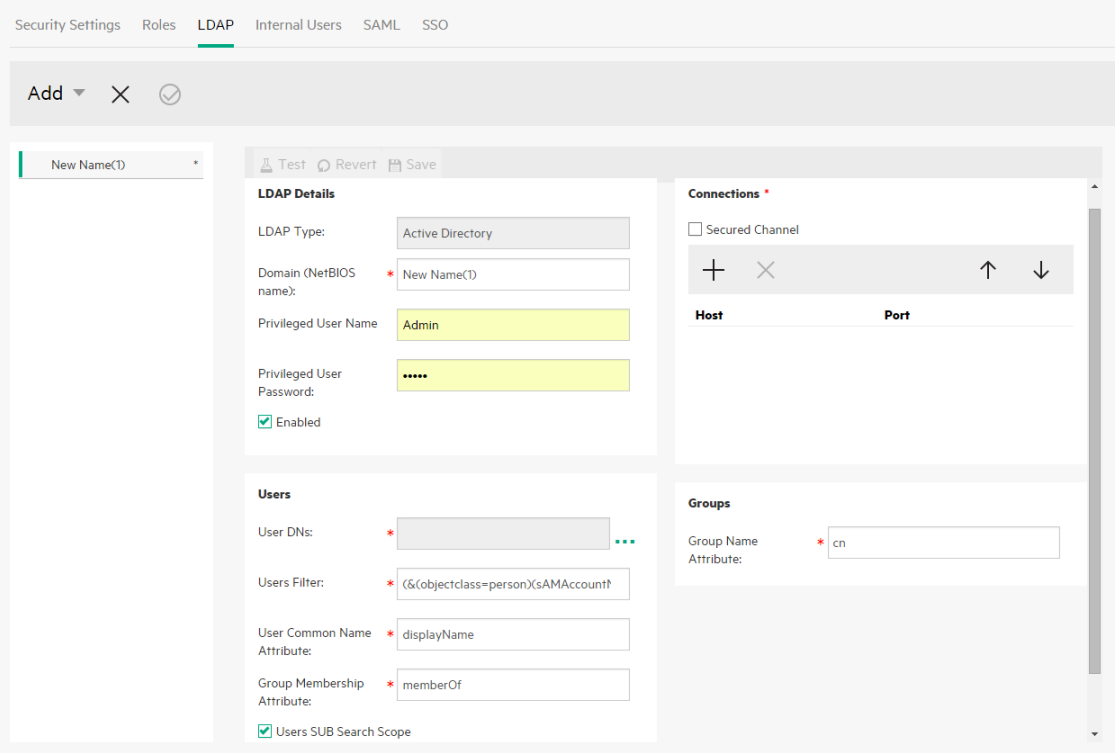

4. In the **Domain (NetBIOS name)** box, enter the AD short domain (NetBIOS name) that you want to authenticate against, for example, EMEA, US, or ASIA. When users log in to HPE OO, they will be able to select from a drop-down list of the active domains.

The domain name appears in the LDAP navigation pane to the left.

5. In the **Privileged User Name** box, enter the user name of a user with search privileges.

   It is recommended to provide a privileged user, because some Operations Orchestration functionality uses searching, for example, the scheduler. If there is no privileged user, this functionality may not work properly.

6. In the **Privileged User Password** box, enter the password of the user with search privileges.
7. Select the **Enabled** check box to enable the authentication of this LDAP.
8. In the **Connections** section, select **Secured Channel** to support secured transport layer connection.
9. In the **Hosts** box, enter the IP address or host name of the LDAP server.
10. In the **Ports** box, enter the port number of the LDAP server. This value must be between 0 and 64435.
11. If you have multiple LDAP servers, click the **Add** + button to add a new line, and add the additional host and port.
12. Use the **Up** and **Down**  arrows to move a connection up and down in the list. A connection that is higher in the list has a higher priority in the case of failover.
13. Click the browse button next to the **User DNs** box, to open the User DNs dialog box, in order to define the distinguished names for users.
    - If you know the user DN, click the **Add** + button to add an empty row, and enter the user DN information manually. If you have multiple LDAP hosts, use ";" as the delimiter.

      For example, `ou=people,dc=devlab,dc=ad;ou=people,dc=devlab,dc=com`
    - If you don't know the exact syntax, click **Add Base DN** to fetch a base DN from the LDAP.

      From the list of available root DNs, select the one that is closes to your needs and then refine

it.

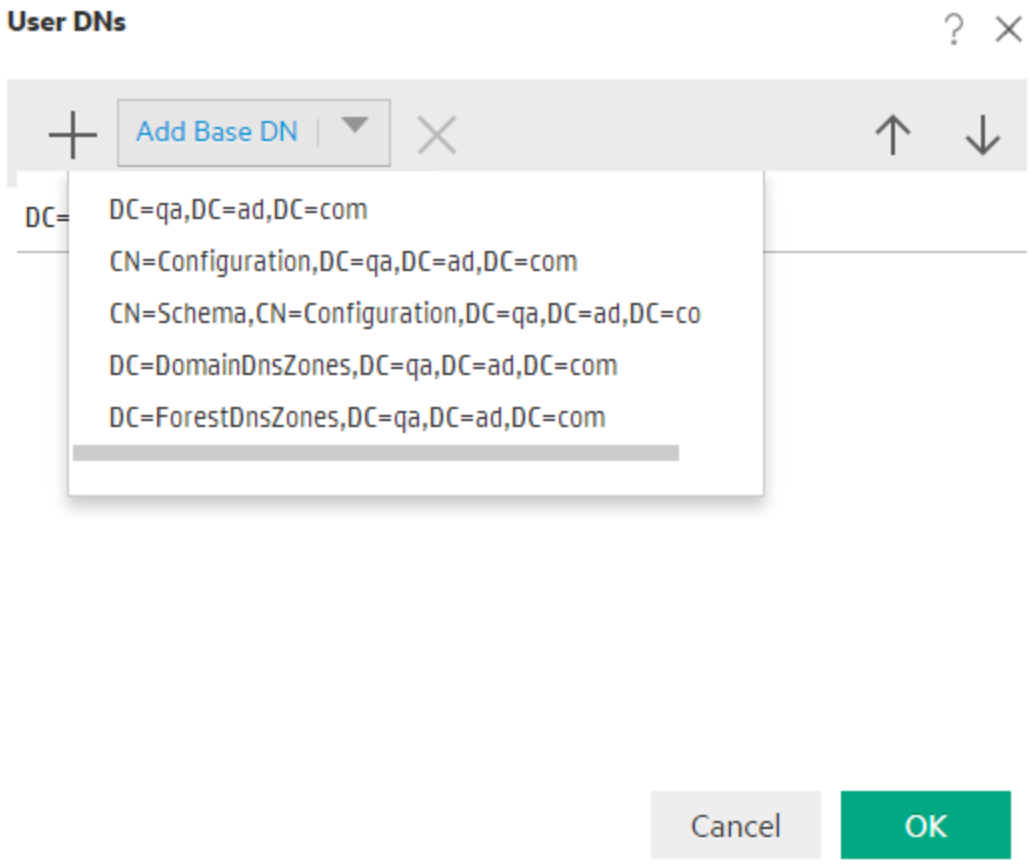

14. Use the **Up** and **Down** arrows  to adjust the order of the user DN rows. The first item in the list is used first.
15. Click **OK** to close the User DNs dialog box and load the information into the **User DNs** box.
16. In the **Users Filter** box, enter the search filter parameters in order to filter users. This filter will be applied on each user DN that was entered in the **Users DNs** box.

    It is recommended to modify the default values, even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute.

    > **Note:** In this filter, `{0}` is a placeholder for the `sAMAccountName` user identifier attribute. When a user logs in, `{0}` is automatically replaced with the supplied user name.

17. In the **User Common Name Attribute** box, enter the attribute that defines how the user name will

be displayed on the screen.

> **Note:** For example, `displayName`.

18. In the **Group Membership Attribute** box, enter the attribute that shows which groups the user belongs to.

    For example, `memberOf`

19. Select the **Users SUB Search Scope** check box if you want the users filter to run recursively.

    > **Note:** If recursive searching is not required, then it is not recommended to select this check box.

20. In the **Group Name Attribute box**, enter the attribute that defines how the name of the group will be displayed in the UI.

    > **Note:** In most LDAP implementations, `cn` is the default.

21. Make sure to test the configuration before saving. For more information, see Testing the Configuration.

22. Click **Save** to save the configuration.

    > **Note:** Make sure to save your changes before enabling authentication. If you don't save, the authentication will not include your changes!

    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.

23. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more information, see "Setting up Security Settings" on page 25.

    > **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.
    >
    > If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

24. In the Enable Authentication dialog box, enter your administrator user name and password, and click **OK**.

Edit an LDAP authentication configuration

1. Select **Security** > **LDAP**.
2. Select the LDAP authentication configuration that you need to edit.
3. Enter the required changes.
4. Make sure to test the configuration before saving. For more information, see Testing the Configuration.
5. Click **Save**.

Configure a list of attributes whose values will be used as groups

It is possible to set up an LDAP authentication configuration with a list of attribute names that will be used as groups, as an alternative to group filters.

This enables organizations to base the functional groups of users on LDAP attributes, rather than LDAP groups.

1. Select **Security** > **LDAP**.
2. From the **Add** menu, select **Alternate Groups**.

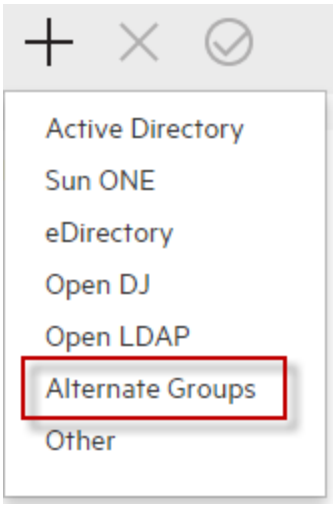

The **LDAP** tab displays the Alternate Groups view. This view is similar to the Non-Active Directory view, but it has an extra text box, **Alternate Group Attributes**, and does not include the **Groups** section.

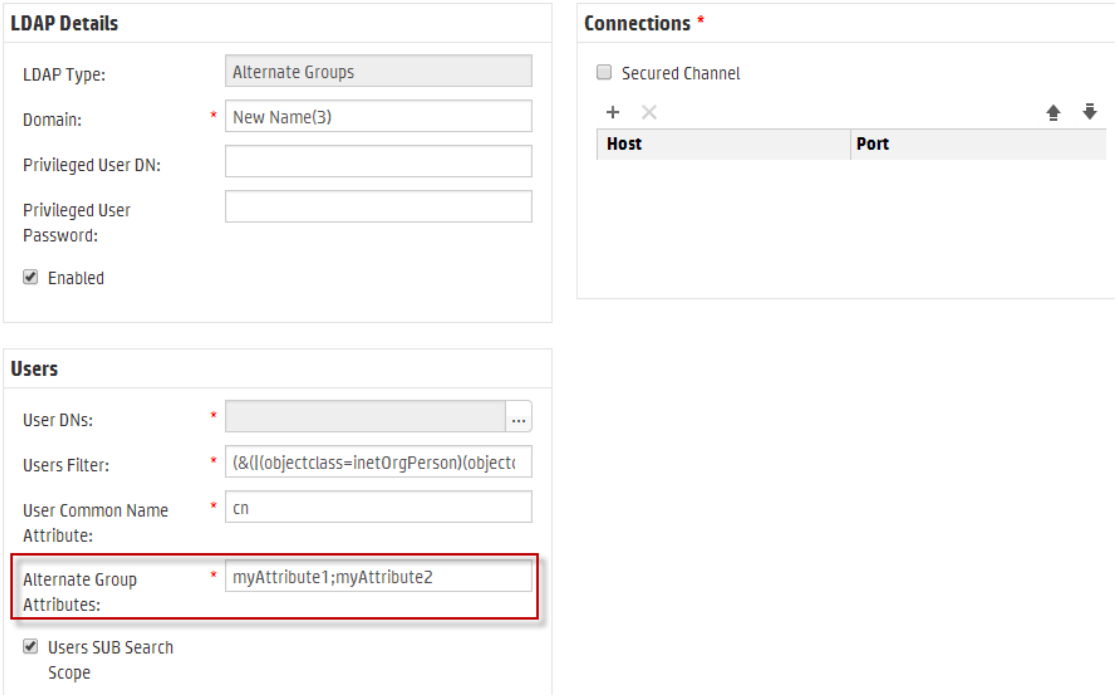

3. In the **Alternate Group Attributes** box, enter the attribute names that will be used as groups. Use a semicolon (;) as the separator.
4. Enter the rest of the configuration details as described in Add an LDAP authentication configuration (not Active Directory), apart from the **Group** fields.

Set the default LDAP authentication configuration

If you set a default LDAP configuration, users who belong to this LDAP will not be required to select a domain, upon login.

Any authentication attempt without a domain will go first to the default LDAP and then to the internal users.

> **Note:** The **Default** button is only available if the selected LDAP configuration has been saved.

1. Select one of the LDAP authentication configurations from the navigation pane to the left.
2. Click the **Default** button in the toolbar.

   The **Default** icon appears next to the configuration name.

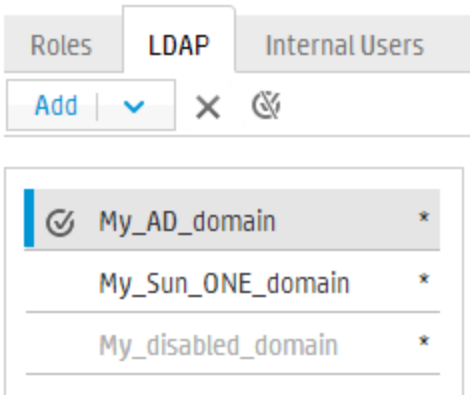

3. To stop this configuration from being the default, select the configuration name and then click the **Remove Default LDAP** button.

Test an LDAP authentication configuration

1. After entering the LDAP details, click the **Test** button.
2. In the Test LDAP Configuration dialog box, enter the user name and password, and click **Test**.

   A pop-up message displays the following information:

   - Whether the authentication is successful
   - Whether the configuration is valid
   - Groups to which the user belongs, if any were found

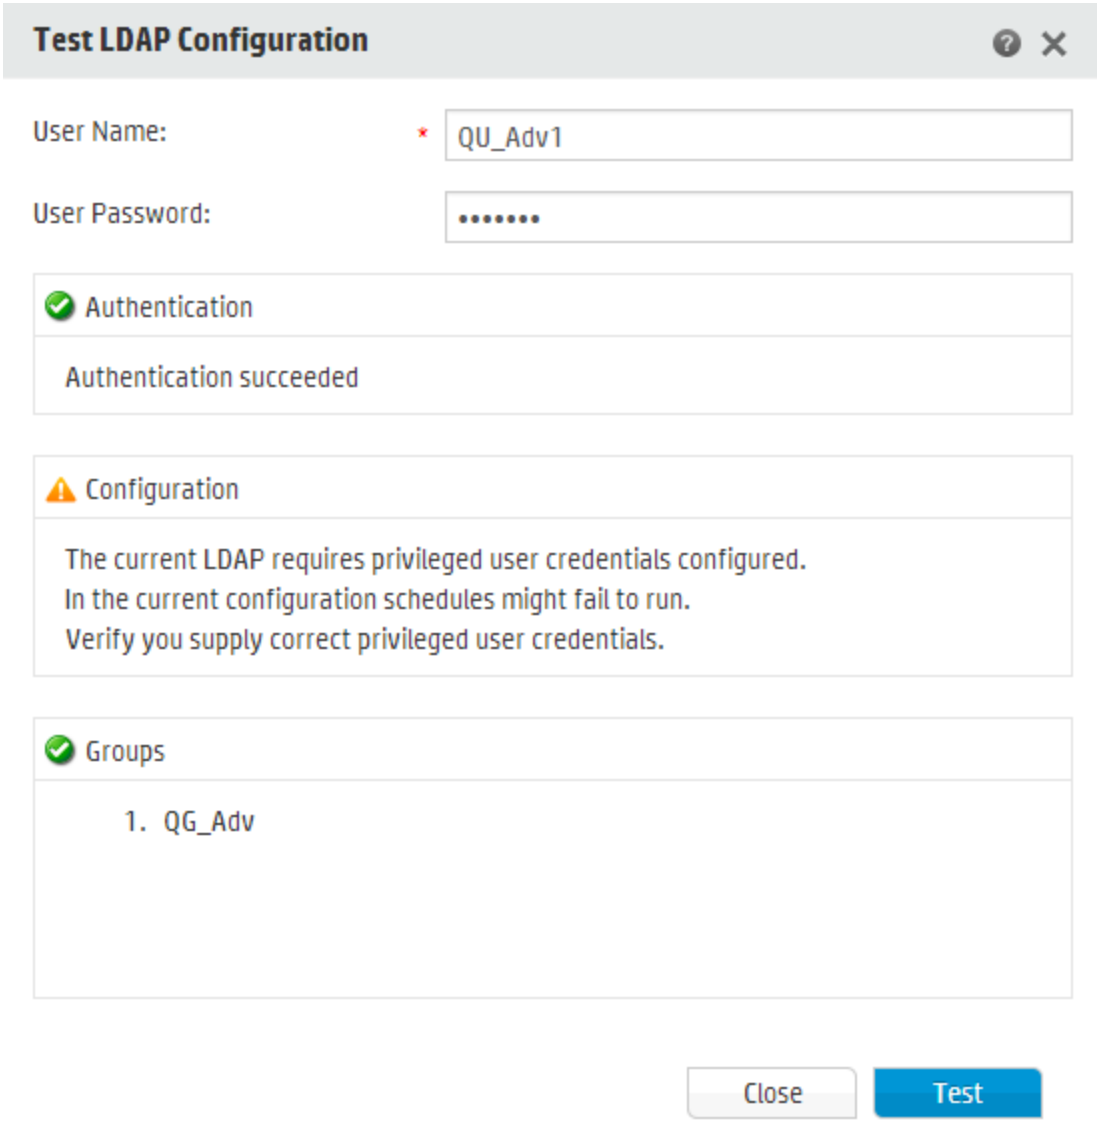

3. Click **Close** to close the dialog box when the testing is complete.

Delete an LDAP authentication configuration

1. Select an LDAP configuration.
2. Click the **Delete** ✕ button in the toolbar.
3. Click **Yes** in the confirmation dialog box.

Disable an LDAP authentication configuration

You might want to temporarily disable an LDAP configuration, while saving the details for later.

1. Select an LDAP configuration.
2. Clear the **Enabled** check box in the **LDAP Details** section.

   The selected LDAP configuration name appears in gray text in the LDAP Navigation pane.

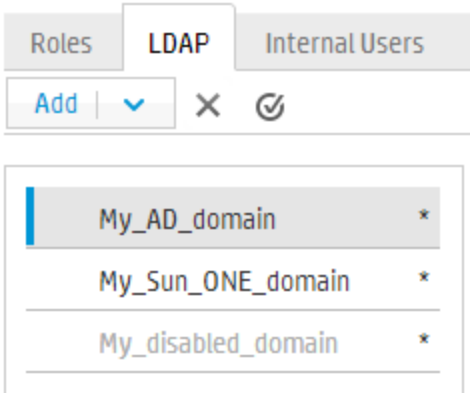

Tell Central whether to trust all valid certificates from secured LDAP servers

When this configuration item is enabled, Central will trust any valid certificate from a secured LDAP server. There will be no need for you to add it to the **client.truststore**.

When this is not enabled, you will need to add the certificates from the LDAP servers used for LDAP authentication to the **client.truststore**. Central will not trust them otherwise.

1. In the command line of the OOSH utility, type (for example):

   ```
   set-sys-config --key ldaps.trustAllCertificates --value true
   ```

2. Press the **Enter** key.

> **Note:** By default, the configuration item named **ldap.trustAllCertificates** is set to **false**. This means that OO will not trust all certificates, and you must configure the LDAP certificate in the Central truststore accordingly.

Reference Material

LDAP Navigation pane

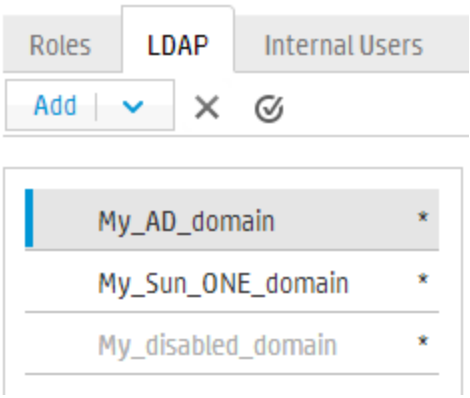

GUI item	Description
Add button	Click to display the list of LDAP types that you can choose from.
Delete button	Click to delete the selected LDAP configuration.
Default button	Click to make the selected LDAP configuration the default. This means that an authentication attempt without a domain will go first to this LDAP and then to the internal users. If you do not set a default, an authentication attempt without a domain will go directly to the internal users. **Note:** The **Default** button is only available if the selected LDAP configuration has been saved.
Navigation pane	Displays the names of the LDAP configuration domain names. Click a domain name to edit the configuration.
Disabled	A domain name appears in gray text if the **Enabled** check box has been cleared for that LDAP configuration.

Security > LDAP>Non-Active Directory configuration

Mandatory fields are indicated with a red asterisk. *

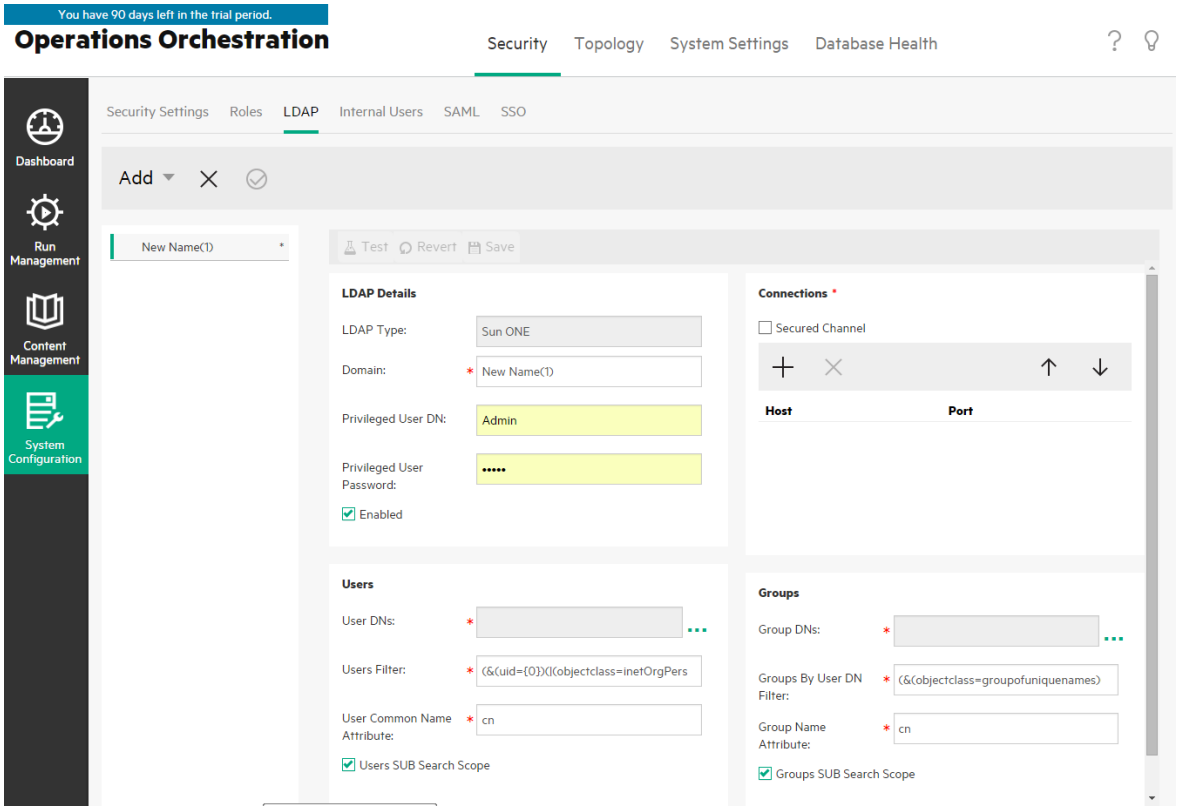

GUI item	Description
Test button	Click to test the configuration. A pop-up message displays the following information: • Whether the authentication is successful • Whether the configuration is valid • Groups to which the user belongs, if any were found
Revert button	Click to revert the configurations to the last saved configurations on the server.
Save button	Click to save the selected LDAP configuration.
LDAP details section	
LDAP Type	Displays the LDAP type that you selected from the **Add** menu.
Domain	Enter a name for the domain. This is name will be displayed when users log in to HPE OO. For example, if you have a single LDAP configuration, you might enter your company name. If you are setting up multiple LDAP configurations, you would give each one a domain that will identify it.

Privileged User DN	Enter the distinguished name of a user with search privileges. For example, `uid=john,ou=people,dc=devlab,dc=ad` 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. **Note:** Make sure to use the exact DN for the user with these privileges.
Privileged User Password	Enter the password of the user with search privileges.
Enabled	Click to enable the authentication of this LDAP
Connections section	
Secured Channel	Select the **Secured Channel** check box to support secured transport layer connection.
Add button	Click to add a new connection.
Delete button	Click to delete the selected connection.
	Use the arrows to move a connection up and down in the list. A connection that is higher in the list has a higher priority in the case of failover
Host	Enter the IP address or host name of the LDAP server. For example, `16.55.222.71`
Port	Enter the port number of the LDAP server. This value must be between 0 and 64435. For example, `389`
Users section	
Users DNs	Define the distinguished names to use for the users. Click the browse button to open the User DNs dialog box. See User DNs/Group DNs dialog box.
Users Filter	Enter the search filter parameters in order to filter users. This filter will be applied on each user DN that was entered in the **User DNs** box. It is recommended to modify the default values, even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute. **Note:** In this filter, `{0}` is a placeholder for the user ID. When a user

	logs in, `{0}` is automatically replaced with the supplied user name.
User Common Name Attribute	Enter the attribute that defines how the user name will be displayed on the screen. **Note:** In most LDAP implementations, `cn` is the default.
Users SUB Search Scope check box	Select this check box if you want the users filter to run recursively. **Note:** If recursive searching is not required, then it is not recommended to select this check box.
Groups section	
Groups DNs	Define the distinguished names to use for the groups. Click the browse button to open the Group DNs dialog box. See User DNs/Group DNs dialog box.
Groups By User DN Filter	Enter the search filter parameters in order to filter the groups. This filters the groups of which the given user is a member. The filter is applied on each Group DN. It is recommended to modify the default values even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute.
Group Name Attribute	Enter the attribute that defines how the name of the group will be displayed in the UI **Note:** In most LDAP implementations, `cn` is the default.
Groups SUB Search Scope check box	Select this check box if you want the groups filter to run recursively. **Note:** If recursive searching is not required, then it is not recommended to select this check box.

Security > LDAP> Active Directory configuration

Mandatory fields are indicated with a red asterisk. *

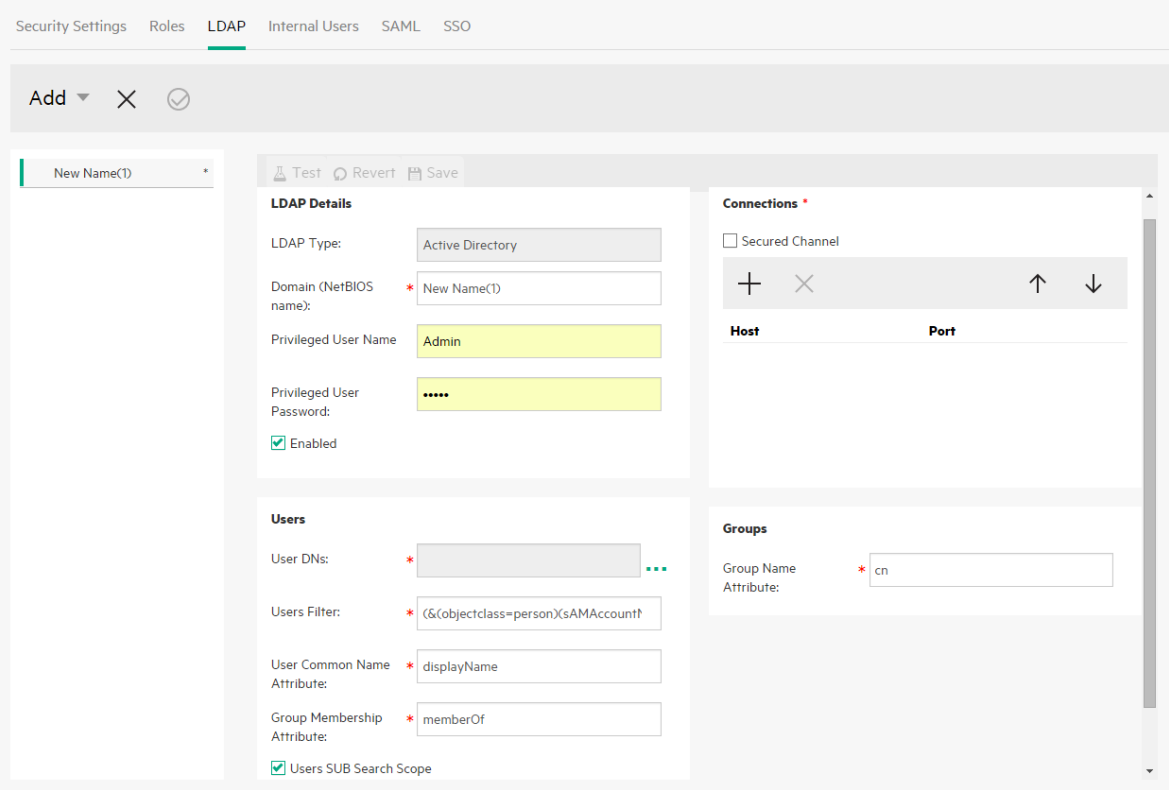

GUI item	Description
Test button	Click to test the configuration. A pop-up message displays the following information: • Whether the authentication is successful • Whether the configuration is valid • Groups to which the user belongs, if any were found
Revert button	Click to revert the configurations to the last saved configurations on the server.
Save button	Click to save the selected LDAP configuration.
LDAP details section	
LDAP Type	Displays the LDAP type that you selected from the **Add** menu.
Domain (NetBIOS name)	Enter the name of the domain (NetBIOS name) that you want to authenticate against. The domain name appears in the navigation pane to the left.
Privileged User Name	Enter the user 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.
Privileged User Password	Enter the password of the user with search privileges.
Enabled	Click to enable the authentication of this LDAP
Connections section	
Secured Channel	Select the **Secured Channel** check box to support secured transport layer connection.
Add button	Click to add a new connection.
Delete button	Click to delete the selected connection.
	Use the arrows to move a connection up and down in the list. A connection that is higher in the list has a higher priority in the case of failover
Host	Enter the IP address or host name of the LDAP server. For example, `16.55.222.71`
Port	Enter the port number of the LDAP server. This value must be between 0 and 64435. For example, `389`
Users section	
Users DNs	Define the distinguished names to use for the users. Click the browse button to open the User DNs dialog box. See User DNs/Group DNs dialog box.
Users Filter	Enter the search filter parameters in order to filter users. This filter will be applied on each user DN that was entered in the **Users DNs** box. It is recommended to modify the default values, even if they work, because this filter is very general. It is highly recommended that the filter contain the `objectclass` attribute. **Note:** In this filter, `{0}` is a placeholder for the `sAMAccountName`user identifier attribute. When a user logs in, `{0}` is automatically replaced with the supplied user name.
User Common Name Attribute	Enter the attribute that defines how the user name will be displayed on the screen.
Group Membership	Enter the attribute that shows which groups the user belongs to.

Attribute	For example, `memberOf`
Users SUB Search Scope check box	Select this check box if you want the users filter to run recursively. **Note:** If recursive searching is not required, then it is not recommended to select this check box.
Groups section	
Group Name Attribute	Enter the attribute that contains the name of the group to be displayed in the UI. **Note:** In most LDAP implementations, `cn` is the default.

Security > LDAP > Alternate Groups configuration

The **Alternate Groups** view is similar to the other non-Active Directory views. However, it does not contain a **Groups** section, and it contains the **Alternate Group Attributes** box.

This enables you to set up an LDAP authentication configuration with a list of attribute names to be used as groups, instead of group filters.

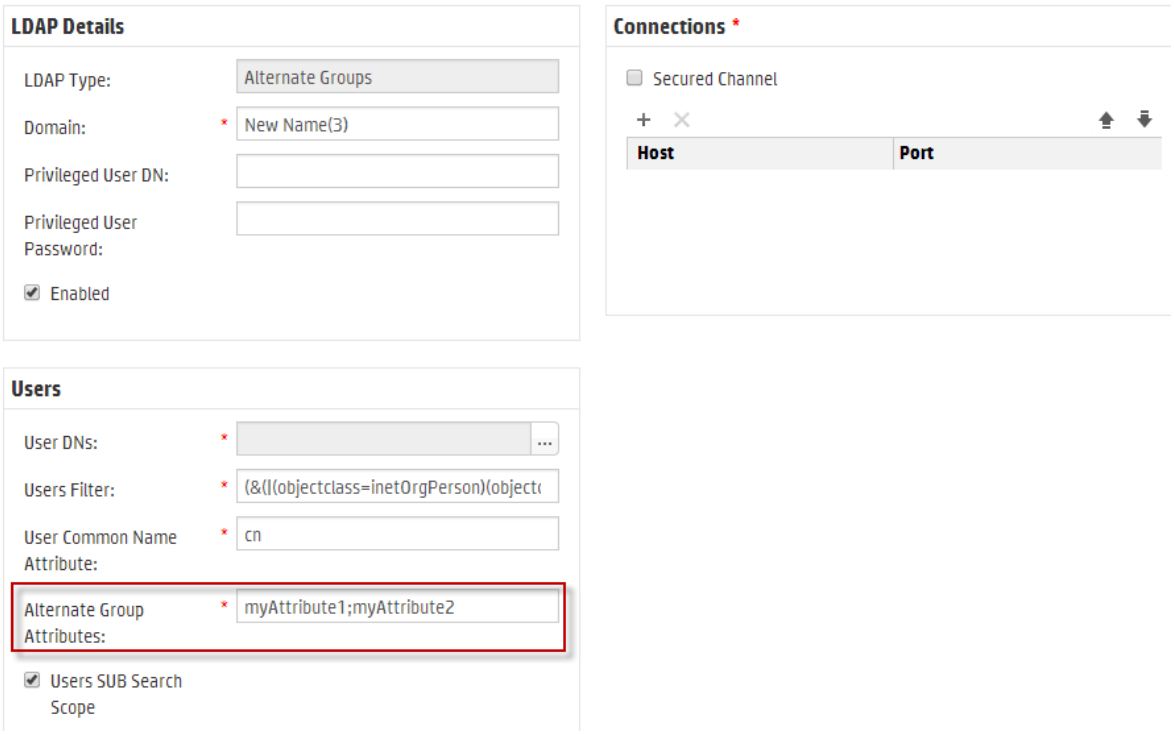

GUI item	Description
Alternate Group Attributes	Enter the attribute names that will be used as groups. Use a semicolon (;) as the separator.

User DNs/Group DNs dialog box

Click the browse button ... next to the **User DNs** box, to open the User DNs dialog box.

Click the browse button ... next to the **Group DNs** box, to open the Group DNs dialog box.

The fields in the two dialog boxes are the same.

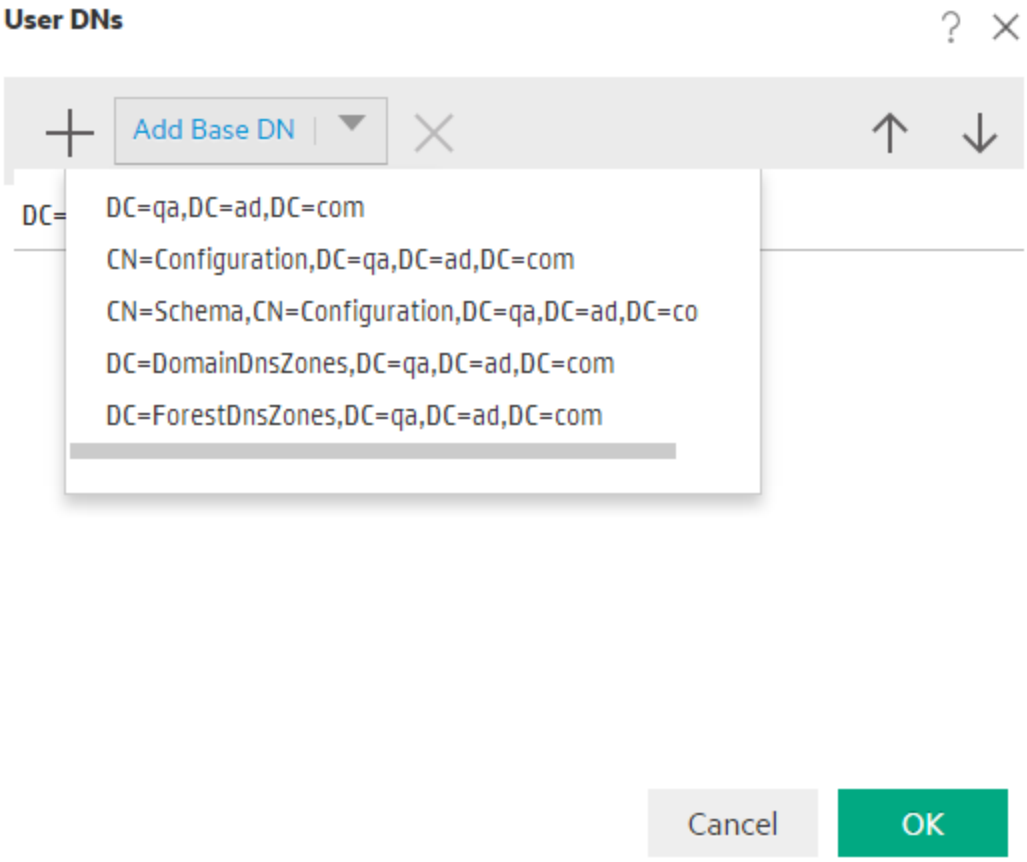

GUI item	Description
Add Base DN	Click to select from the list of available root DNs that are fetched from the LDAP. Select the one that is closes to your needs and then refine it.
Add +	Click to add a new, empty row, in order to manually enter the DN information. If you have multiple LDAP hosts, use ";" as the delimiter. For example, `ou=people,dc=devlab,dc=ad;ou=people,dc=devlab,dc=com`

Delete ✕	Click to delete the selected DN row.
↑ ↓	Use the arrows to adjust the order of the DN rows. The first item in the list is used first.

Troubleshooting

- There is a limit to how many groups a user can have. Users who are members of more than 1,015 groups may fail logon authentication.

  For more information, see http://support.microsoft.com/kb/328889

Setting Up Security - Internal Users

As an alternative to using an external authentication system, such as LDAP, it is possible to set up internal users in Central. Internal users will log in using their HPE OO user name and password, and they will not require any external authentication.

By assigning roles to users, the administrator controls which actions users are allowed to perform in the system. For more information about roles, see "Setting Up Security – Roles" on page 32.

> **Note:** The **Internal Users** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the users if you have been assigned a role with the **Manage Security Configuration** permission.

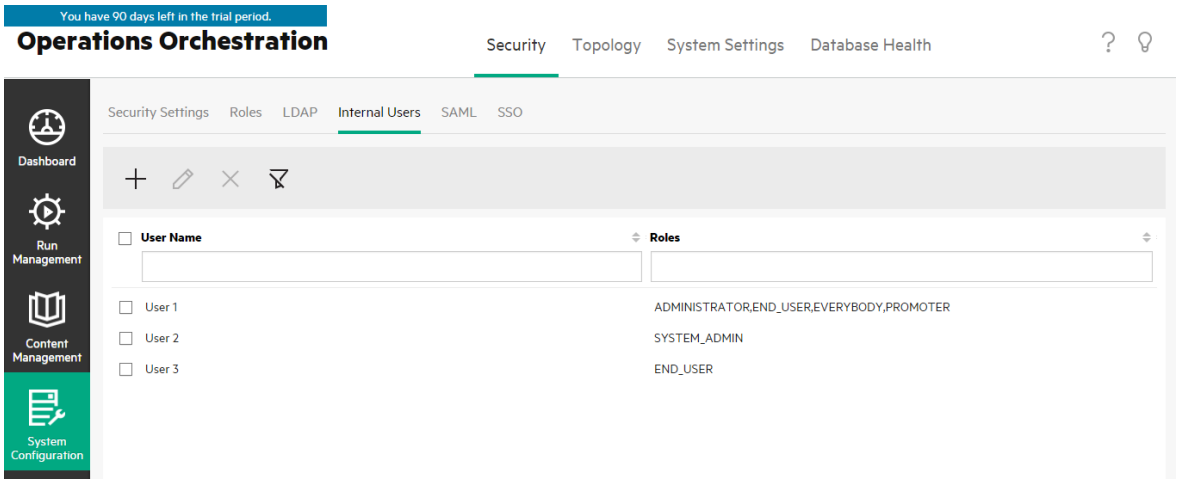

If a user has been assigned the same user name for both LDAP user and internal user, Central will relate to the LDAP user.

What do you want to do?

Add a user

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **Internal Users**.
3. Click the **Add** button.
4. Enter the details for the internal user:

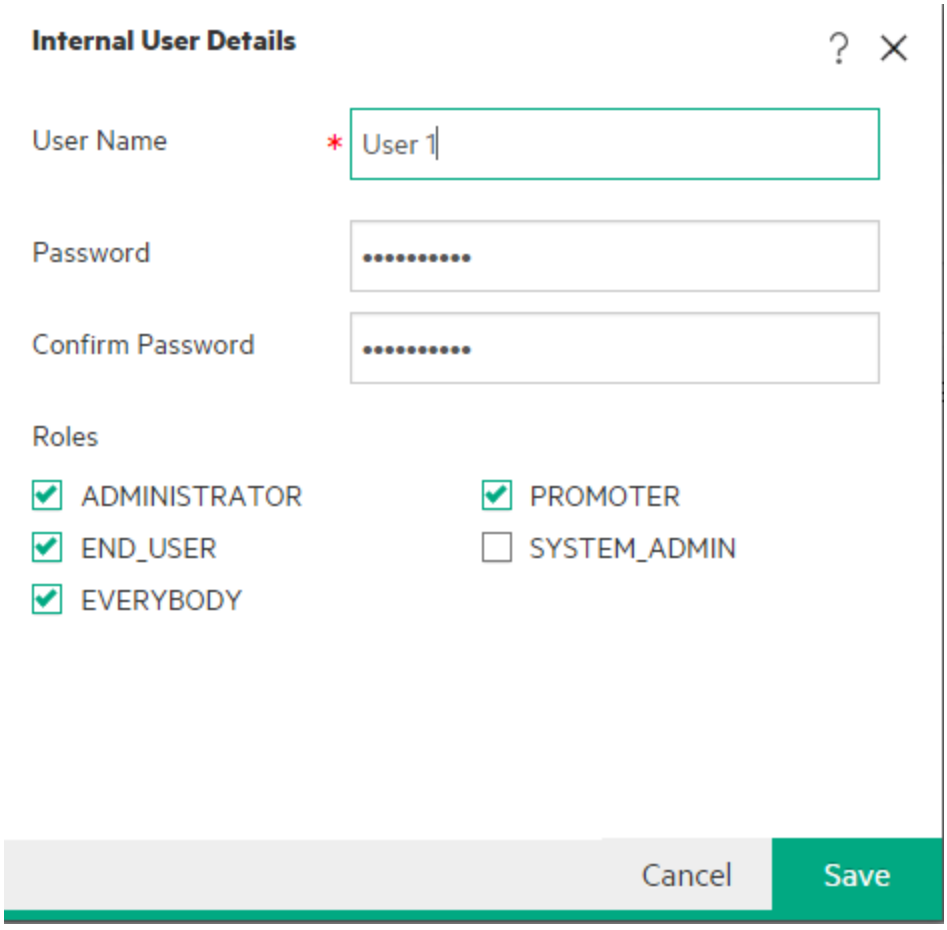

a. In the **User Name** box, enter a unique user name. The user name cannot have more than 255 characters.

b. Enter the password in both the **Password** and the **Confirm Password** boxes. The password cannot have more than 255 or less than 6 characters.

c. Select the check box next to each role that you want to assign to this user. You can assign multiple roles to a user.

> **Note:** If you assign multiple roles to a user, the user will be able to access all the permissions in those assigned roles. A role with fewer permissions will not restrict this user's access to the permissions in their other roles.

5. Click **Save** to save the new user.
6. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more information, see "Setting up Security Settings" on page 25.

> **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.
>
> If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

Edit a user account

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **Internal Users**.
3. Select the check box next to the user that you want to edit and click the **Edit**  button.

> **Note:** The **Edit** button is only enabled if a single user is selected.

4. Make the required changes in the User Configuration dialog box.
5. Click **Save**.

Delete a user

1. Select **Security** > **Internal Users**.
2. Select one or more users that you want to delete and click the **Delete** ✕ button.

> **Note:** The **Delete** button is only enabled if at least one user is selected.

3. Click **OK** in the confirmation dialog box.

Note: A user cannot delete their own account.

Reference Material

Security > Internal Users

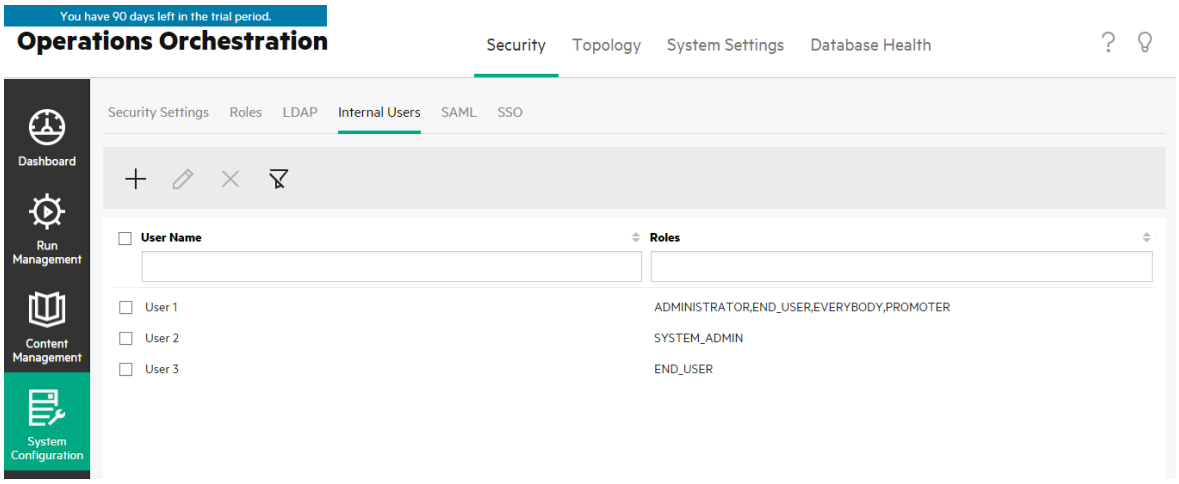

GUI item	Description
Add button	Opens the Internal Users Configuration dialog box, so that you can add a new user.
Edit button	Opens the Internal Users Configuration dialog box, so that you can edit the selected user.
Delete button	Deletes the selected user.
User Name	Displays the user name.
Roles	Displays the roles that have been assigned to the user.
Filter boxes	In the filter boxes at the top of the columns, enter text to filter the displayed users.

Internal User Details dialog box

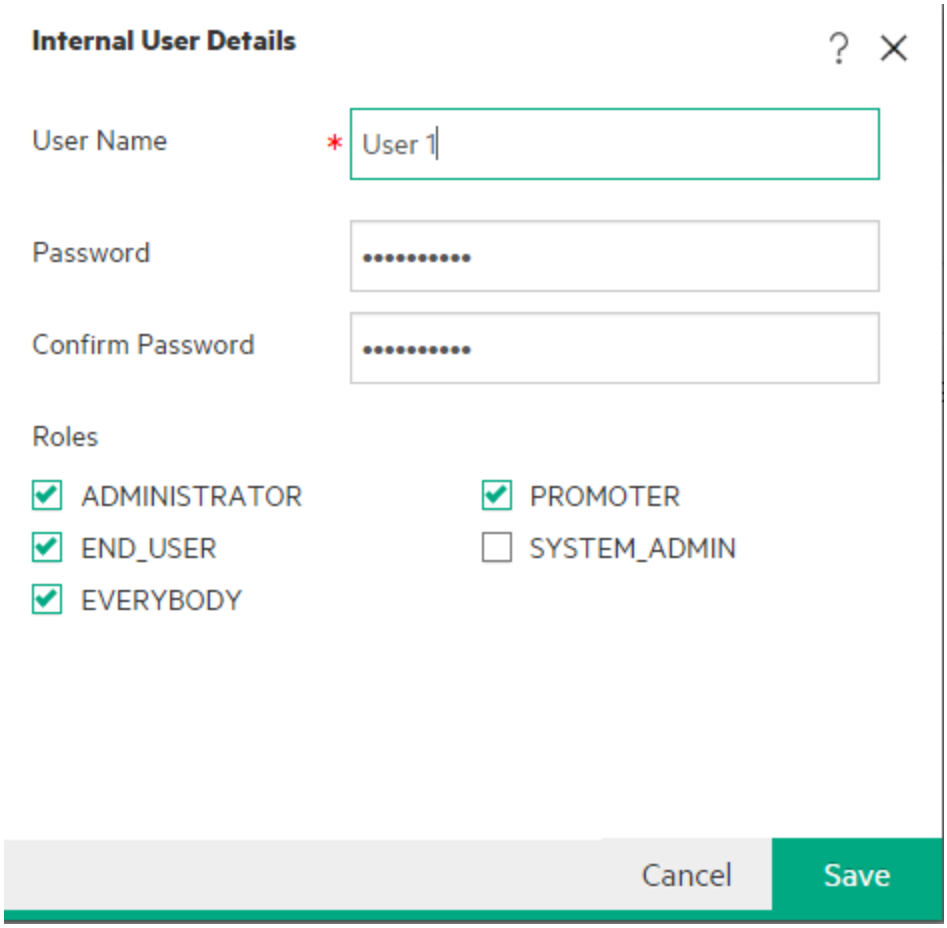

GUI item	Description
User Name	Enter a unique user name of up to 255 characters.
Password	Enter a password of 6 to 255 characters.
Confirm Password	Enter the password a second time for verification.
Roles	Select the check box next to each role that you want to assign to this user. You can assign multiple roles to a user.
Save	Click to save the user.

Setting Up Security – SAML Authentication

 Administrators can configure Central to work with SAML 2.0 authentication. SAML (Security Assertion Markup Language) is an XML-based open standard data format for exchanging authentication and authorization data between an identity provider and a service provider. The process of authenticating and authorizing users with the identity provider involves the use of a SAML artifact and SAML assertions.

A **SAML assertion** is an XML document that includes the attributes that define the user and the user's group. The identity provider sends these attributes to the service provider, and based on these attributes, the service provider provides access rights to the user.

A **SAML artifact** is a sequence of bytes that encodes two pieces of information:

- The artifact ID for the SAML assertion, which is retrievable from the identity provider
- A SOAP endpoint, which can be used to resolve this handle.

The SAML artifact string should be passed as a query parameter over the request URL in Base64 encoding format.

For example:

```
http(s)://<FQDN>:<PORT>/saml/SSO/alias/defaultAlias?SAMLart=ARTIFACTEncodedString
```

> **Note:** The **SAML** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the SAML authentication if you have a role with the **Manage Security Configuration** permission.

> **Tip:** It is recommended to keep the admin user, so that you will be able to log in directly, if there is an error in your SAML configuration.

> **Note:** You must use an **internal/ldap** user when using the HPE OO Shell Utility (OOSH), the Studio remote debugger, when creating a flow schedule in Central, and when using RESTful APIs or backwards compatible SOAP and RESTful APIs from HPE OO 9.x.

Known Issues and Troubleshooting

- If you get locked out, with no access via the identity provider, log into HPE OO directly. Make sure to keep the admin user, for this. To log in directly, go to **http(s)/<host>:<port>/oo/login/direct**.
- The Central Scheduler does not work when SAML authentication is used.

> **Note:** If you are using LDAP, you must set a default LDAP configuration. See "Set the default LDAP authentication configuration" in "Setting Up Security – LDAP Authentication" on page 42.

What do you want to do?

Prerequisites

Before setting up SAML authentication, complete the following steps:

1. Configure HPE OO for SSL. See the *HPE OOSecurity and Hardening Guide*.
2. If you have changed the **key.store** default password, you need to set it as a system configuration item.

   For instructions, see Change the default values of the SAML system configuration. For reference information, see SAML system configuration.

   > **Note:** This step is not needed if you use the default password.

3. If you have changed the server certificate (private key) default password and the alias default value of this certificate in the **key.store**, you need to set them as system configuration items.

   For instructions, see Change the default values of the SAML system configuration. For reference information, see SAML system configuration.

   > **Note:** This step is not needed if you use the default values.

4. Import the public keys of the identity provider's certificate into the Central **key.store**.

   For example:

   ```
   keytool -importcert -alias <any_alias> -keystore key.store -file <certificate_
   name.cer> –storepass <changeit>
   ```

5. Restart the server.
6. Open the Central user interface.

   Important! Make sure to use the fully qualified domain name in the URL.

7. Create the administrator user and administrator role, and any other roles that are needed.
8. Map the roles to the required groups. For more information about creating roles and mapping them to groups, see "Setting Up Security – Roles" on page 32.

Change the default values of the SAML system configuration

If you have changed the **key.store** default password, you need to set it as a system configuration item.

If you have changed the server certificate (private key) default password and the alias default value of this certificate in the **key.store**, you need to set them as system configuration items.

> **Note:** These steps are not needed if you use the default values.

1. Connect to the HPE OO Shell utility (OOSH).

   For more information about theOperations Orchestration Shell utility, see the Operations Orchestration *Shell User Guide*

2. If you want to change the password of the **key.store** or the Operations Orchestration server certificate (that was imported into the **key.store**), you need to encrypt the password with the encryption tool, located at **<install dir>/central/bin/encrypt-password**:

   ```
   encrypt-password --encrypt --password <your password>
   ```

   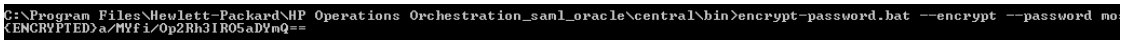

3. Run `set-sys-config --key <key> --value <value>`

   For example:

   ```
   set-sys-config --key key.store.private.key.alias.name --value newValue
   ```

   > **Note:** You can see a list of all system configurations by running the '`lsc`' command.

See SAML system configuration in the **Reference** section.

Add a SAML authentication configuration

Before starting this procedure, read the Prerequisites section, and complete any necessary steps.

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **SAML**.
3. In the **Service Provider** section, enter the entity ID of the service provider (HPE OO). The default is **ooentityid**.

In the **Service Provider** section, enter the entity ID of the service provider (HPE OO). The default is **ooentityid**.

4. In the **Identity Provider** section, enter the details of the identity provider:
   - **IDP metadata URL**: Enter the SAML metadata for the identity provider with which you will interact.
   - **User name attribute**: Enter the SAML assertion attribute name that defines the user.
   - **Group names attribute**: Enter the SAML assertion attribute name that defines the user groups.
   - **Group names delimiter**: Enter the delimiter used for the groups in the SAML assertion group names attribute.

   > **Note:** Make sure that the attributes are correct; otherwise, access will be denied.
5. If HPE OO is installed behind a proxy, define the forward (web) proxy host and port number.
6. Click **Enable** to enable SAML.
7. Click **Save**.
8. Click the **Download Metadata** button to save the metadata file.

   > **Note:** This button is enabled after you enable and save the SAML configuration.
9. Provide the HPE OO public key to the identity provider.
10. Provide the HPE OO public key to the identity provider.
11. Provide the metadata file to the identity provider.
12. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more information, see "Setting up Security Settings" on page 25.

    > **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.
    >
    > If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

Add SAML authentication configuration to a load balancer

When setting up a load balancer, you need to resend the service provider metadata to the identity provider.

Before starting this procedure, read the Prerequisites section, and complete any necessary steps.

1. Complete all the steps in the Add a SAML authentication configuration section.
2. Set the external URL of the load balancer (under the **Topology** tab) and save. For more information, see "Setting Up Topology – Workers and RASes" on page 81.
3. Click the **Download Metadata** button to get the new metadata with the external URL details.
4. Provide the metadata file to the identity provider.

Change the SAML log level

You can change the SAML log level in the **log4j.properties** file.

1. Open the **log4j.properties** file (under **/<oo-installation>/central/conf/log4j.properties**).
2. Locate the line that refers to the SAML logging level. For example:

   ```
   saml.log.level=INFO
   ```

3. Change the logging level to `DEBUG` or `ERROR/WARNING` as required.

For more information, see "Adjusting the Logging Levels" in the *HPE OO Admistration Guide*.

For more information, see "Adjusting the Logging Levels" in the *HPE OO Admistration Guide*.

Reference Material

Security > SAML

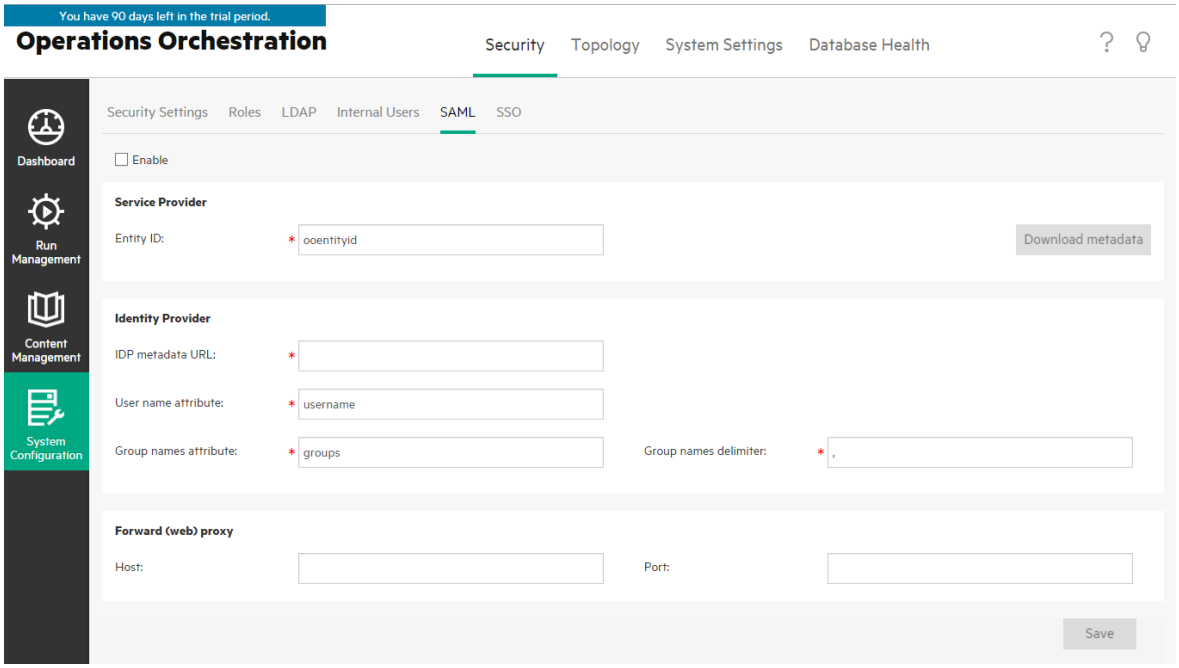

GUI item	Description
Enable check box	Select to enable the SAML authentication
Entity ID	Enter the identity ID of the service provider (HPE OO). The default is **ooentityid**.
IDP metadata URL	Enter the SAML metadata for the identity provider with which you will interact.
User name attribute	Enter the SAML assertion attribute name that defines the user.
Group names attribute	Enter the SAML assertion attribute name that defines the user groups.
Group names delimiter	Enter the delimiter used for the groups in the SAML assertion group names attribute.
Host	If HPE OO is installed behind a proxy, define the forward (web) proxy host.

Port	If HPE OO is installed behind a proxy, define the forward (web) proxy port number.
Download metadata	Click to save the metadata file, which then needs to be sent to the identity provider. This button is enabled after you enable and save the SAML configuration.

SAML system configuration

If you update the following configuration items using the HPE OO Shell utility (OOSH), you will need to restart the server (security context). For more information about the HPE OO Shell utility, see the HPE OO *Shell User Guide*.

Configuration item	Action
`key.store.password`	To set the password used to access to the **key.store**. The default value is 'changeit'.
`key.store.private.key.alias.name`	To set the alias used to for the server certificate (private key) in the **key.store**. The default value is 'tomcat'.
`key.store.private.key.alias.password`	To set the password used to access the server certificate (private key) from the **key.store**. The default value is 'changeit'.

Setting Up Security - LWSSO

 You can set up single sign-on (SSO) for HPE OO through Lightweight SSO (LWSSO). LWSSO is an HPE solution that enables single sign-on using one authentication across various HPE applications.

LWSSO shares a cookie between HPE products that are accessed from a web browser. As a result, if a user logs onto another HPE product web client that has LWSSO enabled, such as the SM web client or the BSM web client, this user will be able to enter the HPE OO Central application directly, bypassing the HPE OO Central logon screen.

Prerequisites for this configuration:

- The LWSSO feature must be enabled on all the HPE products between which you want to use single sign-on.
- The user's logon credentials in the other HPE product must match those of an HPE OOuser account.

  For example, if BSM is integrated with HPE OO, the user that is logged into BSM must also exist inHPE OO (either an LDAP user or an internal user).

> **Note:** Even if HPE OO is set up with multiple LDAP configurations, it is only possible to authenticate the user using LWSSO with the default LDAP. LWSSO will first try to authenticate the user with the default LDAP, and if this fails, will try to authenticate within the HPE OO internal domain.

After you save an LWSSO configuration, the changes take effect immediately, and there is no need to restart the server.

All nodes in the cluster configuration will refresh their configuration from the database periodically at a 10 minute interval. In the case of a single server, the node will update its configuration from the database at the same 10 minute interval.

When connecting another application that employs LWSSO to HPE OO, you must specify the connection URL of HPE OO using the following format:

```
<protocol>://<FQDN>:<portNumber>
```

For example, `http://lab.lab:8080`

> **Note:** The LWSSO enabling procedure may be different for other HPE products. See the corresponding documentation for each HPE product.

> For example, for some applications, the connection URL needs to be in the format `<protocol>://<FQDN>:<portNumber>/oo`.

> **Note:** The **SSO** tab is only visible if you have been assigned a role with the **Manage Security Configuration** or **View Security Configuration** permission. You will only be able to edit the LWSSO authentication if you have a role with the **Manage Security Configuration** permission.

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, in the **ras/conf/ras-wrapper.conf** file. For more information, see "Configuring LWSSO Settings" in the *HPE OO Administration Guide*.

> **Note:** If you are using LDAP, you must set a default LDAP configuration. See "Set the default LDAP authentication configuration" in "Setting Up Security – LDAP Authentication" on page 42.

What do you want to do?

Set up the LWSSO authentication configuration

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Security** > **SSO**.

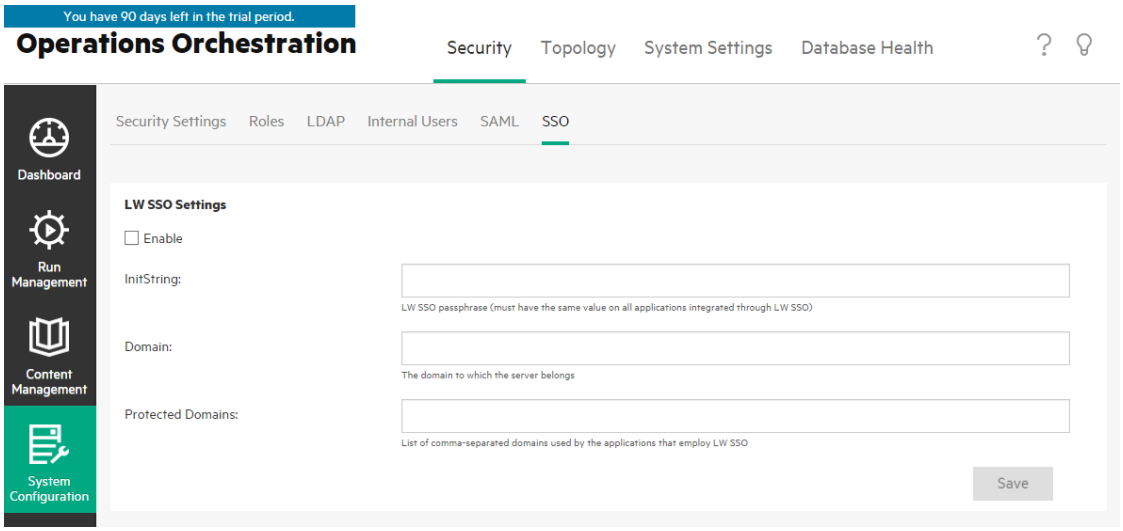

3. In the **initString** box, enter the password you want to use to connect HPE products. This value must be the same as those used in the LWSSO configuration of the other HPE products. This password must contain at least 12 characters.

The initstring is used to encrypt the LWSSO cookie and has no link to the user's password. However, when two applications are integrated using LWSSO, the initstring in both applications must match. Otherwise, the cookie encrypted by one application cannot be decrypted by the other application.

4. In the **Domain** box, enter the domain of the HPE OO server.
5. If there are multiple domains used by applications that employ LWSSO, enter them in the **Protected Domains** box, using commas as separators.

   The HPE OO server domain must appear in the **Protected Domains** list. When there are several domains in the **Protected Domains** box, all applications that are integrated through LWSSO must define the same list of protected domains.
6. Select the **Enable** check box to enable the LWSSO configuration. If you do this, after you save the settings, the icon next to the **LW SSO Settings** title will change to **Enabled**  .

   > **Note:** The **Enable** check box is there to let you disable an LWSSO configuration while keeping the general authentication feature enabled. For example, you might have authentication enabled for internal users and LDAP users, but you may or may not want the LWSSO configuration to be enabled.
   >
   > Note that the new state of LWSSO settings is not enabled immediately when you select the **Enable** check box, but when you save the new configuration.
7. Click **Save**.
8. To enable the authentication functionality, select the **Settings** tab (under **Security**) and select the **Enable Authentication** check box to display the Enable Authentication dialog box. For more information, see "Setting up Security Settings" on page 25.

   > **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.
   >
   > If the **Enable Authentication** check box is not selected, there is no authentication and all users will be able to access all tasks.

Edit an existing LWSSO authentication configuration

1. Select **Security** > **SSO**.
2. Enter the required changes.
3. Click **Save**.

Reference Material

Security > SSO

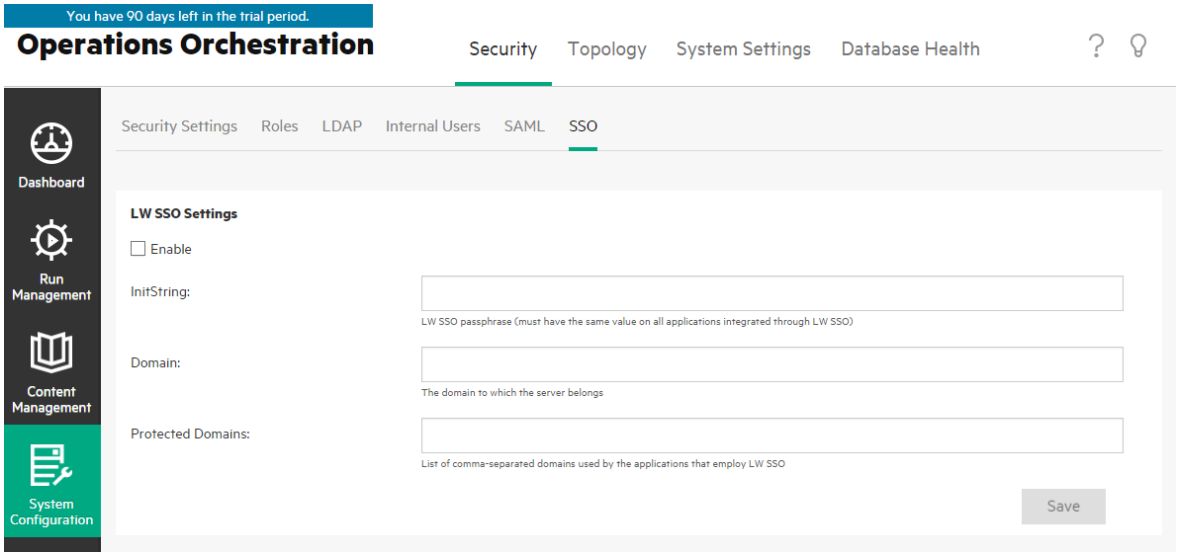

GUI item	Description
Enable	Select the check box to enable LWSSO.
initString	The initstring is used to encrypt the LW SSO cookie and has no link to the user's password. However, when two applications are integrated using LWSSO, the initstring in both applications must match. Otherwise, the cookie encrypted by one application cannot be decrypted by the other application.
Domain	Enter the domain of the HPE OO server.
Protected Domains	If there are multiple domains used by applications that employ LWSSO, enter them in the **Protected Domains** box, using commas as separators. The HPE OO server domain must appear in the **Protected Domains** list. When there are several domains in the **Protected Domains** box, all applications that are integrated through LWSSO must define the same list of protected domains.
Save	Click to save the LWSSO authentication configuration.

Setting Up Topology – Workers and RASes

 Network topology is the arrangement of the various elements (links, nodes, and so on) in a network. Workers and RASes are part of the topology and can be configured under the **Topology** tab.

The **Workers** tab displays information about the following items:

- **Central worker** - Central workers are responsible for executing flows. A Central worker obtains tasks (flow execution messages) to process from Central.

  A Central worker can be assigned to a worker group. This enables workers to withstand a high action execution load and increases availability of workers in a data center. See Assigning Workers to Groups.

- **Standard RAS** - a RAS is a remote action server, containing a worker and a remote protocol for connecting with Central. A standard RAS will initiate the connection to Central.

  After a standard RAS is installed, it will appear in the grid in the **Workers** tab.

- **Reverse RAS** - a reverse RAS will wait for Central to initiate the connection. To define a reverse RAS, you must:

  a. Install the reverse RAS during the installation procedure. This creates the **ras-connectivity.properties** file containing the RAS configuration details. See "Configure Central to register the Reverse RAS" in the *Installation, Upgrade, and Configuration Guide*.

  b. Then, register the reverse RAS in the **Workers** tab. See Registering a Reverse RAS.

  > **Note:** One of the Centrals is selected to establish connection to the RAS. If this Central is shut down, all RASes that this Central was responsible to establish connection to, will be reassigned to other Central nodes in the cluster equally.

  While attempting to connect to a reverse RAS:

  - If the RAS does not answer after 30 seconds, a timeout will occur.
  - If there are several consecutive failed connection attempts (because the wrong shared secret was entered), this will result in a delay. If this occurs, check that you are using the correct shared secret.

> **Note:** The **Topology** tab is only visible if you have been assigned a role with the **Manage**

> **Topology** or **View Topology** permission. You will only be able to edit the topology if you have a role with the **Manage Topology** permission.

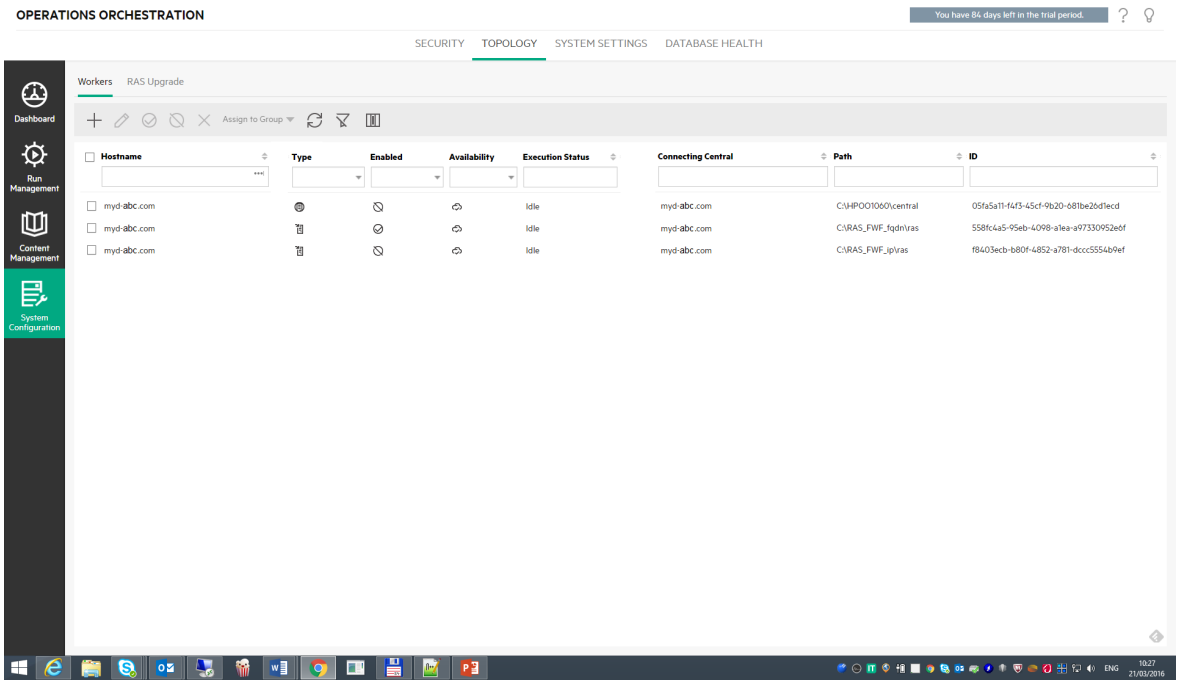

The **Topology** > **Workers** tab displays information about each worker/RAS: their status (**Running** or **Stopped**), installation path, connection details, operating system, ID, and so on.

> **Note:** The status can take a minute or two to update. If you stop a RAS service, it takes a few minutes until this is reflected in the **Topology** > **Workers** tab. There is no need to restart Central.

Execution Status

The **Topology** > **Workers** tab displays information about the resources that are being consumed by each worker/RAS, under **Execution Status**. This information may help with troubleshooting the system, and may contribute to decisions about scaling and resource allocation.

The information about resource consumption is automatically refreshed periodically. Take into account that the displayed information is only relevant to the last few minutes, since the previous refresh, and is not cumulative. In order to properly understand what is happening in your environment, it is recommended to observe the statuses over a period of time.

A worker/RAS can have one of the following execution statuses:

- **Idle**: This worker has almost no tasks to process.
- **Low Utilization**: This worker is working at low utilization.

If your workers are with idle/low utilization over a long period of time, you might consider removing them.

- **Working**: This worker is working with a regular load.
- **Pending on Central**: Most of the time, this worker is waiting to send execution results to Central. Central is busy with handling other tasks.

  If this state is displayed over a long period of time, you might consider adding more Centrals, or monitoring your database for space, connectivity, and so on, to make sure it is not a bottle neck.
- **Loaded**: At the current time, this worker works at maximum ability and is not able to get more tasks from Central. If all your workers are loaded most of time, consider adding more Centrals or RASes.
- **N/A**: No information is provided for this worker.

Registering a Reverse RAS

While installing a RAS, you chose whether the RAS will initiate the connection to Central (Standard RAS) or wait for Central to initiate the connection (Reverse RAS)

If a RAS was configured to accept connection from Central (Reverse RAS), you must configure Central to register the RAS by providing all required information: host/IP, port, and so on. Until you do this, the RAS will be idle, waiting for Central to initiate the connection.

> **Note:** If you configured the RAS to initiate connection with Central (Standard RAS), then there is no need to register the RAS.

> **Note:** You cannot connect a reverse RAS to more than one Central.

Assigning Central Workers to Groups

A worker group is a logical collection of Central workers. Having a group, rather than a single worker, enables workers to withstand a high action execution load and increases availability of workers in a data center.

> **Note:** A worker may belong to more than one group simultaneously.

Under the **Topology** > **Workers** tab, you can assign workers to worker groups, remove workers from groups, enable, disable, or delete workers.

> **Note:** Some execution stages always run on the default group, which is **RAS_Operator_Path**. Therefore, make sure that at least one worker is assigned to this group.

It is recommended to assign most (or even all) workers to **RAS_Operator_Path**, in order to optimize the execution time. You should only exclude a RAS from this group if you want to minimize its usage, and if you have only few operations that use it.

For more information about workers and worker groups, see "Worker Groups and Group Aliases" in the *HPE OO Concepts Guide*.

What do you want to do?

Register a new reverse RAS

For each reverse RAS that you installed (those that were configured to accept connection from Central), you must configure Central to register the RAS. Until you do this, the RAS will be idle, waiting for Central to initiate the connection.

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **Topology** > **Workers**.
3. Click  **Add Reverse RAS**.
4. In the dialog box, enter the RAS connection details: host name or IP, port, and protocol.

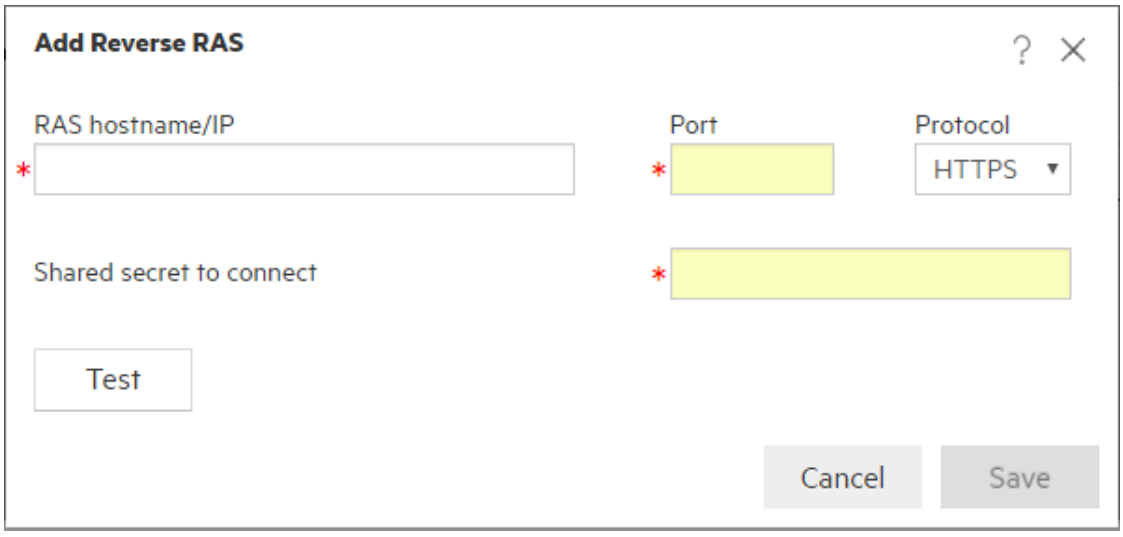

5. Enter the shared secret, which was entered when the reverse RAS was installed.
6. Click **Test**.

   If you encounter an error, adjust the details and try again.
7. Click **Save**.

Central initializes the connection and the connection remains active. If the connection becomes disconnected, Central will try to reconnect at specified intervals.

The new reverse RAS appears in the table below.

Move a reverse RAS to a different port or IP address

1. Select **Topology** > **Workers**.
2. Select the reverse RAS you want to move.
3. Click  .
4. In the dialog box, change the reverse RAS connection details: host name or IP, port, and protocol.

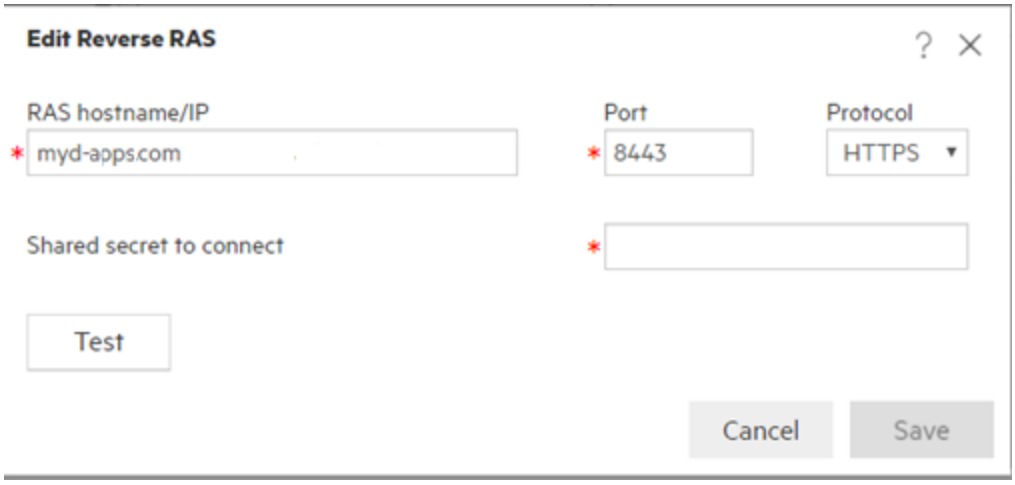

5. Click **Test**.

   If you encounter an error, adjust the details and try again.
6. Click **Save**.

The new reverse RAS appears in the table below.

Assign a Central worker to an existing worker group

1. Select **Topology** > **Workers**.
2. Select the check box next to the worker name.

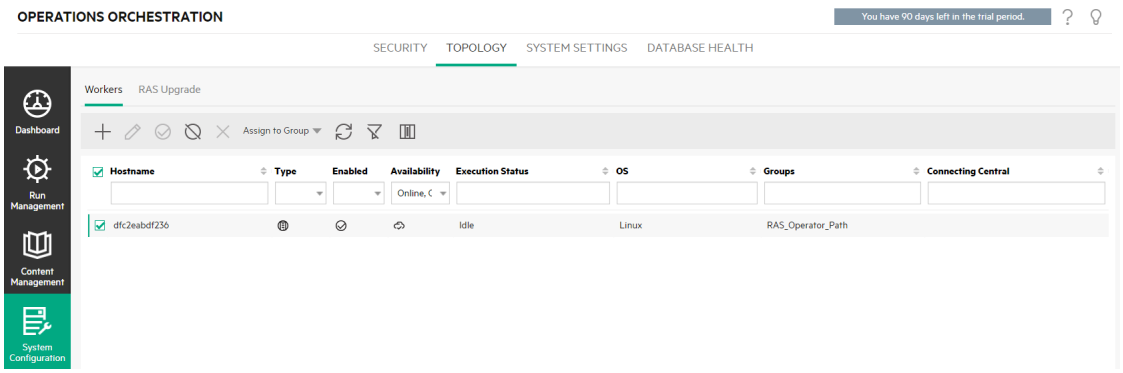

3. Click the **Assign to Group** button.
4. Select the check boxes next to the worker groups that you want to assign this worker to.

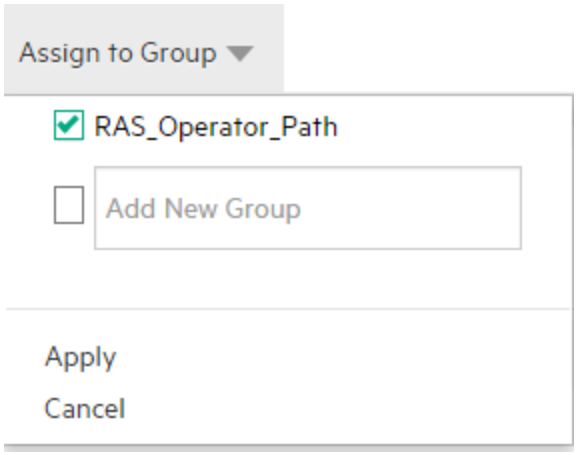

5. Click **Apply**.

Create a new worker group and assign a Central worker to it

1. Select **Topology** > **Workers**.
2. Select the check box next to the worker name.
3. Click the **Assign to Group** button.
4. Click the check box next to the empty box and enter the name of the new worker group that you want to create.

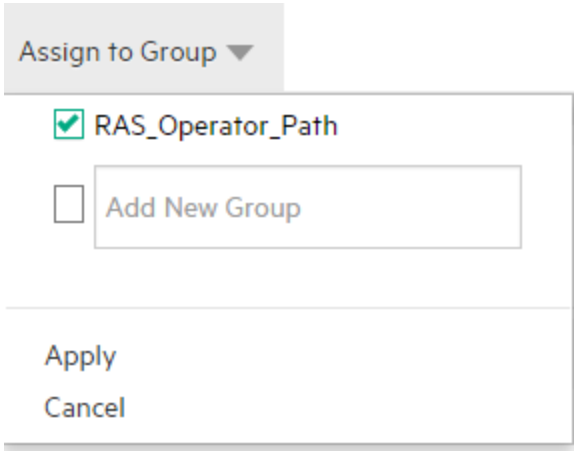

5. Click **Apply**.

Remove a Central worker from the a worker group

1. Select **Topology** > **Workers**.
2. Select the check box next to the worker name.
3. Click the **Assign to Group** button.
4. Clear the check boxes next to the worker groups that you want to remove this worker from.
5. Click **Apply**.

Disable a worker/RAS

1. Select **Content Management** > **Topology** > **Workers**.
2. Select the check box next to the worker name.
3. In the Workers toolbar, click **Disable** .

> **Note:** The **Disable** button is only available if at least one enabled worker is selected.

Enable a worker/RAS

1. Select **Topology** > **Workers**.
2. Select the check box next to the name of the disabled worker.
3. In the Workers toolbar, click **Enable** .

> **Note:** The **Enable** button is only available if at least one disabled worker is selected.

Delete a RAS

You can delete a RAS only if it is stopped.

It is not possible to delete an internal worker, which is within Central.

> **Note:** If you have stopped a RAS service (or even uninstalled the RAS), it takes a few minutes until this is reflected in the **Topology** > **Workers** tab. You may need to wait for a few minutes before that status is updated to **Stopped**, before you can delete the RAS. There is no need to restart Central.

1. Select **Topology** > **Workers**.
2. Select the check box next to the worker name.
3. In the Workers toolbar, click **Delete**  .

> **Note:** The **Delete** button is only available if at least one worker is selected and if the selected worker is not currently running.

Verify that a component (RAS or cluster node) was successfully installed

You can use the **Topology**/**Workers** tab to verify that a component was successfully installed.

If a worker is successfully installed, it will be visible in the grid and shown to be online.

Reference Material

Topology > Workers

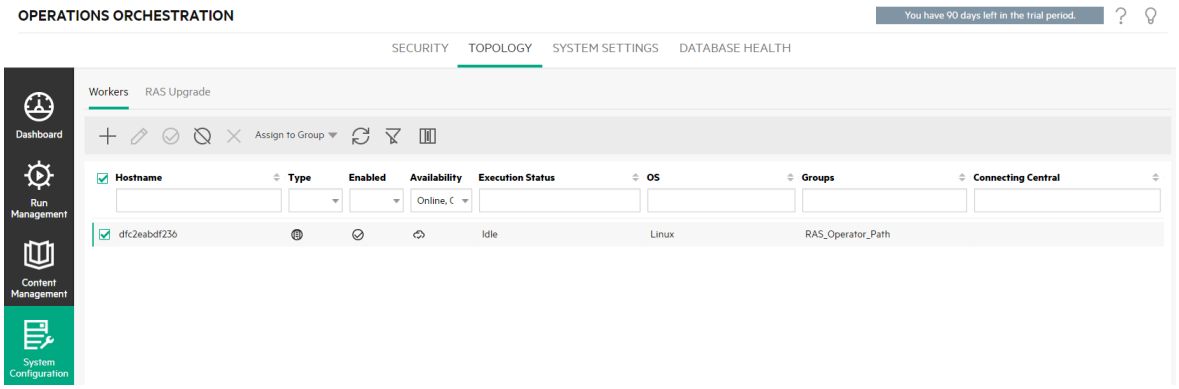

GUI item	Description
Add Reverse RAS button 	Click to create a new reverse RAS connection. You must do this for each reverse RAS that you installed (those that were configured to accept connection from Central).
Edit Reverse RAS button 	Click to modify the connection details of the selected reverse RAS.
Enable button 	Click to enable the selected worker/RAS. The Enable button is only available if at least one disabled worker/RAS is selected.
Disable button 	Click to disable the selected worker/RAS. The Disable button is only available if at least one enabled worker/RAS is selected.
Delete button 	Click to delete the selected worker/RAS. The Delete button is only available if at least one worker/RAS is selected and if the selected worker/RAS is not currently running.
Refresh button 	Click to refresh the display of workers/RASes.
Clear Filters button 	Click to remove the filters and display all the workers/RASes.
Select Columns button 	Click to display the column picker, to choose which columns to display.
Assign to Group button	Click to display the **Assign to Group** pop-up menu.
Filter boxes	In the filter boxes at the top of the columns, enter text to filter the displayed workers/RASes.
Host Name	The name of the host where this worker/RAS is located.
Enabled	Whether the worker/RAS is enabled or disabled.
Availability	Whether the worker/RAS is online or offline.
Execution Status	Displays information about the resource consumption of the worker/RAS. The information about resource consumption is automatically refreshed periodically. Take into account that the displayed information is only relevant to the last few minutes, since the previous refresh, and is not cumulative. In order to properly understand what is happening in your environment, it is recommended to observe the statuses over a period of time. • **Idle**: This worker/RAS has almost no tasks to process.

<table>
<tr><td></td><td>

Low Utilization: This worker/RAS is working at low utilization. If your workers/RASes are with idle/low utilization for a long while, consider removing them.
Working: This worker/RAS is working with a regular load.
Pending on Central: Most of the time, this worker/RAS is waiting to send execution results to Central. Central is busy with handling other tasks. Consider adding more Centrals, or monitor your database for space, connectivity, and so on, to make sure it is not the bottle neck.
Loaded: At the current time, this worker/RAS is not able to process all Central tasks. If a worker/RAS is in a loaded state for a long time, consider scaling out or scaling up by adding more worker threads.

</td></tr>
<tr><td>Path</td><td>The path to the host where this worker is located. The path indicates if the worker is in Central or in a RAS.</td></tr>
<tr><td>OS</td><td>The operating system of the machine where this worker/RAS is located.</td></tr>
<tr><td>JDK</td><td>The JDK version that is installed on the worker/RAS machine.</td></tr>
<tr><td>.NET</td><td>The .NET version that is installed on the worker/RAS machine.</td></tr>
<tr><td>ID</td><td>The UUID (unique identifier) of the worker/RAS.</td></tr>
<tr><td>Groups</td><td>The worker groups to which the selected worker/RAS is assigned.</td></tr>
<tr><td>Type</td><td>Whether the item is a Central worker, a standard RAS, or a reverse RAS.
<table>
<tr><td>Central worker  </td><td>Central workers are responsible for executing flows. An external worker connects to Central to obtain tasks (flow execution messages) to process. It can be assigned to a worker group. See Assigning Workers to Groups.</td></tr>
<tr><td>Standard RAS  </td><td>A RAS is a remote action server, containing a worker and a remote protocol for connecting with Central. A standard RAS will initiate the connection to Central.</td></tr>
<tr><td>Reverse RAS  </td><td>A reverse RAS will wait for Central to initiate the connection. See Registering a Reverse RAS.</td></tr>
</table>
</td></tr>
<tr><td>Connecting Central</td><td>For a reverse RAS only: The host name of the Central that connects to the reverse RAS.</td></tr>
</table>

Connect Reverse RAS dialog box

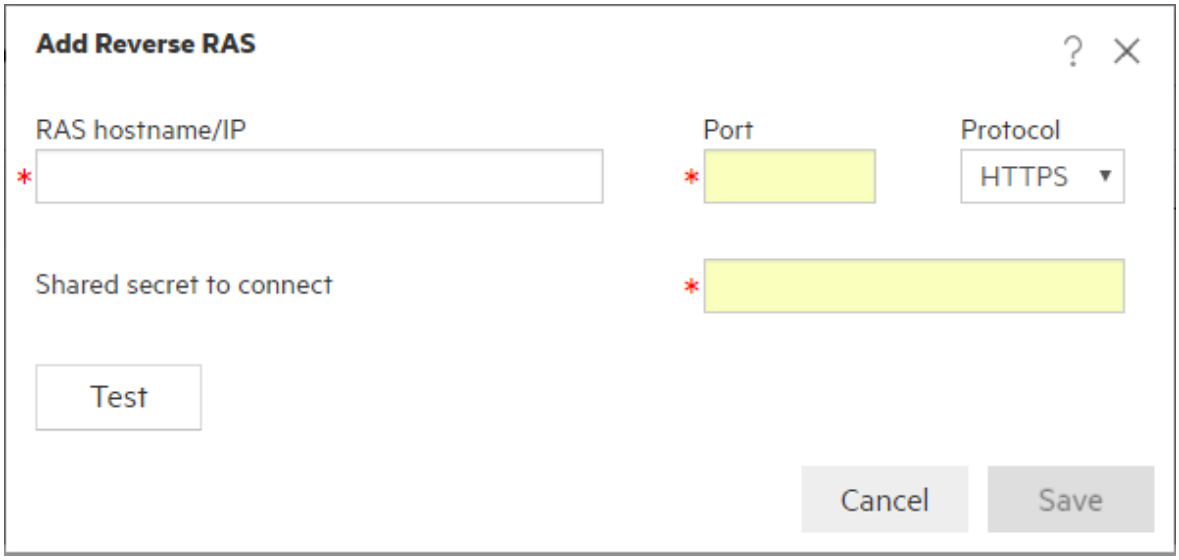

GUI item	Description
RAS host name/IP	Enter the name or IP of the RAS (mandatory).
Port	Enter the port of the RAS (mandatory).
Protocol	Enter whether the protocol is HTTP or HTTPS.
Share secret to connect	Enter the shared secret that was entered in the installation wizard when this RAS was installed (mandatory).
Test	It is recommended to test the connection before submitting.

Topology > Workers > Assign to Group pop-up

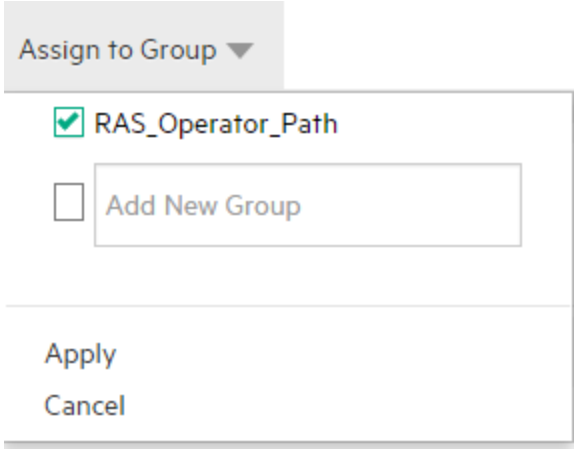

GUI item	Description
Assign to Group	Click the button to display the **Assign to Group** pop-up menu.
Add New Group	Select the check box and enter the name of a new group, to which the selected worker will be assigned.
<Worker groups>	Select the check boxes next to group names to assign the selected worker to these groups. The display in the dialog box will vary, depending on the groups that have been set up.
Apply	Click to apply the changes to the worker.
Cancel	Click to close the Assign to Group pop-up without saving the changes.

Setting Up Topology – Automatic RAS Upgrade

 Your Centrals and RASes must be the same version. Central will only assign jobs to RASes that are of the same version.

Since version 10.6x, HPE OO supports the automatic upgrade of RASes from Central, without manual intervention and without needing to access the RAS machines physically.

Automatic RAS upgrade is useful for customers with a large number of RASes in their environment, or with RASes in different time zones, making manual upgrading difficult.

> **Note:** For FIPS140-2 compliant RASes you cannot use the automatic RAS upgrade procedure, so the RASes must be upgraded manually. This is because for FIPS140-2 compliance, the OpenJDK needs to be replaced by the Oracle JDK. In addition, the JDK needs to be configured for FIPS compliance which causes the upgrade package's signature to be voided.
>
> As the upgrade package must have a valid signature in order for the automatic RAS upgrade procedure to be successful, this procedure cannot be used.

> **Note:** Automatic RAS upgrade is only supported when both Central and the RASes are already at version 10.60 or later. For earlier versions, you will still need to upgrade the RASes manually. For example, you can't use it to upgrade RASes from 10.5x to 10.60.

> **Note:** The **RAS Upgrade** tab is only visible if you have been assigned a role with the **Manage Topology** or **View Topology** permission.

There are four steps that you need to complete in order to upgrade your RASes using Automatic RAS Upgrade:

1. Upload the **RAS-Upgrade.zip** file to Central. This is a zip file containing the files that are required to upgrade the RASes to the relevant version. You can download the **RAS-Upgrade.zip** file from the same location where you download the new OO installation. Then, you upload the file to the Central database.
2. Prepare the RASes for upgrade. In this step, the RAS automatically runs a set of preparations, to ensure that the upgrade will be successful. This step runs pre-checks to verify that there are no problems in each RAS environment, and prepares those environments for upgrade. The RASes

then download the upgrade package from Central.

a. To start the preparation process, you need to manually select the required RASes from the grid and click the **Prepare for Upgrade**  button.

   During the preparation process, each RAS displays its status in the RAS grid. At the end, it displays whether the process was successful.

   If a RAS encounters a problem, the details are displayed on the RAS grid. You need to fix the problem before upgrading that RAS. Examples of problems might be:

   - Insufficient disk space. The RAS upgrade file may require up to 2 GB.
   - Missing OS permissions (administrator permissions).
   - Another HPE OO component (Central or Studio) is installed alongside the RAS in the same HPE OO installation directory. This kind of setup is not supported by this feature, and will require a manual upgrade.
   - Problems where you need to check in the logs for more information.

b. Once you have fixed the problems, you should run the preparation process again. This process can be repeated multiple times.

During the first two steps, you can use the RASes as usual. There is no need to schedule downtime.

The **RAS-Upgrade.zip** file is signed and verified before uploading and also after downloading. If packages are lost, the download will fail, so if this occurs, the download should be restarted.

When RAS transfers the **RAS-Upgrade.zip** file, this is done by streaming very small packets. This means that it can be sent to multiple RASes simultaneously, while taking up only minimal resources. This also means that if the download fails, on retry, it will start from where it failed, so the packets that were transferred successfully don't need to be transferred again.

3. Upgrade Central. The next step is to upgrade your environment, by upgrading Central to the new version. This is done manually. For more information, see the *HPE OO Installation, Upgrade, and Configuration Guide*.
4. Upgrade the RASes. After Central is upgraded, you need to upgrade your RASes to the same version.

   Important! The versions of Central and the RASes ***must be aligned***. RASes with older versions will not be able to run flows.

   In the RAS grid, the **Version** column displays the version of each RAS. Check this column to see if your RASes are aligned with the Central version.

> **Note:** A RAS can only be upgraded if the preparation was successful, and its **Upgrade State** shows as **Ready for upgrade**.

To start the upgrade process, you need to manually select the required RASes and click the **Upgrade Now**  button.

During this process, each RAS runs the upgrade batch file, which is located in the **Upgrade-RAS.zip** file.

When the process ends, the **Upgrade State** column in the RAS grid shows whether the upgrade was successful. Each RAS that was successfully upgraded restarts and reports to Central that it is up in a new version.

The first two steps should be performed before the Central upgrade. This minimizes downtime and enables you to fix any issues that may be encountered, before the upgrade.

> **Note:** If Central is rolled back to the previous version, you will have to manually roll back the upgraded RASes.

> **Note:** This feature only works if you are upgrading from 10.6x to a later version of 10.6x. It will not work for RASes from earlier versions.

For more information about RASes, see the *Operations Orchestration Concepts Guide*.

What do you want to do?

Step 1: Upload the upgrade file to Central

This preparatory step should be completed before Central is upgraded.

1. Get a local copy of the **RAS-Upgrade.zip** file for the new version. You can do one of the following:
   - Go to the new HPE OO version page in the HPE SSO portal, and download the **RAS-Upgrade.zip** file and save it to your local file system. (This is supported for versions later than 10.60 only).
   - Go to the new HPE OO version page in the HPE SSO portal, and download the complete Central upgrade/installation package for the new version, and save it to your local file system.
2. Click the **System Configuration** button.
3. Select **Topology** > **RAS Upgrade**.

4. Click the **Upload** button.
5. In the Upload File dialog box, click **Browse** to navigate to the **RAS-Upgrade.zip** file, in the location where you saved it.

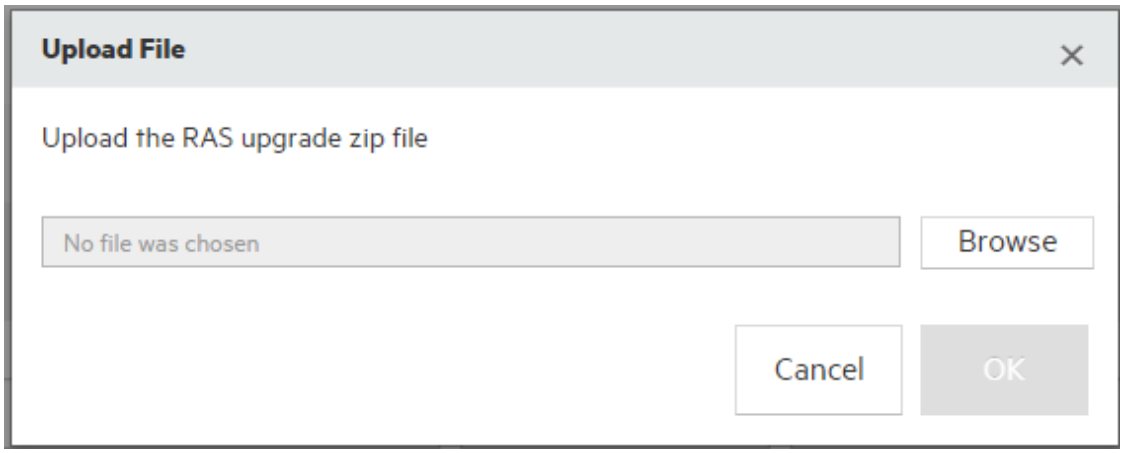

> **Note:** Make sure to upload the **RAS-Upgrade.zip** file, and not the full upgrade package.

6. Click **OK** to upload the upgrade file to Central.

   Central performs validations (for example, verifies the digital signature of the upgrade file), and receives the upgrade file.

Step 2: Prepare the RASes for upgrade

This step prepares the environment for upgrade by running pre-checks on the RASes, checking for disk space, permissions, and so on, and downloading the upgrader package to the RASes. The preparation is done in the background, and doesn't interrupt the routine work of your RASes.

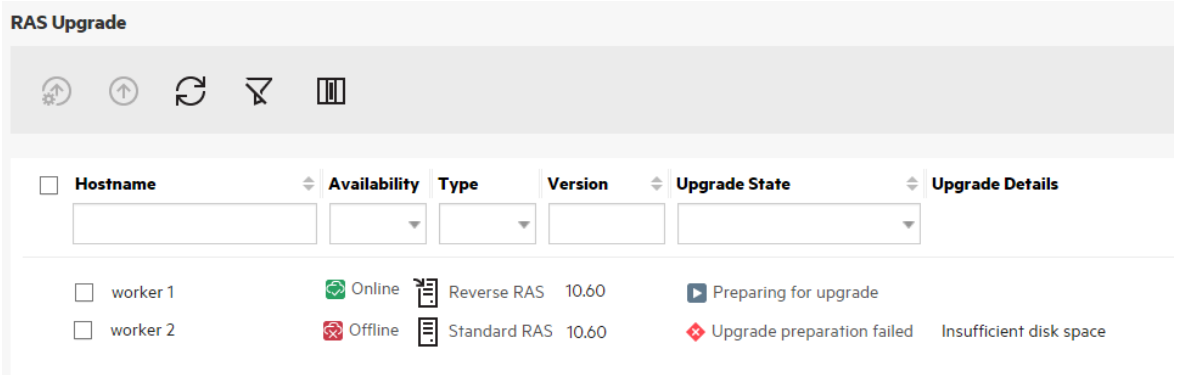

1. In the **Topology** > **RAS Upgrade** tab, select one or more RASes and click the **Prepare for Upgrade** button.

2. Check the status of the RASes, as displayed in the **Upgrade State** column.

   If there are any issues encountered during the download of the upgrader package to any of the RASes, the details are displayed next to that RAS.

3. If required, fix the problems and click the **Prepare for Upgrade**  button again.

   When the checks and download are completed successfully for a RAS, the **Upgrade State** for that RAS is updated to **Ready for Upgrade**.

> **Note:** You can prepare RASes for upgrade either before or after you upgrade Central.
>
> You can only prepare RASes for upgrade if you have uploaded the **RAS-Upgrade.zip** file to Central.

Step 3: Upgrade Central

Upgrade Central to the relevant version, by running the upgrade script.

For more information, see the *HPE OO Installation, Upgrade, and Configuration Guide*

Step 4: Upgrade the RASes

This step should be done after Central has been upgraded and after the RASes have been prepared for upgrade.

> **Note:** A RAS can only be upgraded if the preparation was successful, and its **Upgrade State** shows as **Ready for upgrade**.

1. Select **Topology** > **RAS Upgrade**.
2. Select the RASes that you want to upgrade (from those that passed the preparation), and click the **Upgrade Now**  button. This triggers the auto upgrade in the RASes.

   > **Note:** When you click the **Upgrade Now**  button, the RASes perform the prerequisite checks again. This is done because time may have passed since when you clicked the **Prepare for Upgrade**  button, and the situation may have changed.

3. Follow the **Upgrade State** to monitor when your RASes have the status of **Upgrade succeeded**.

Reference Material

Topology > RAS Upgrade

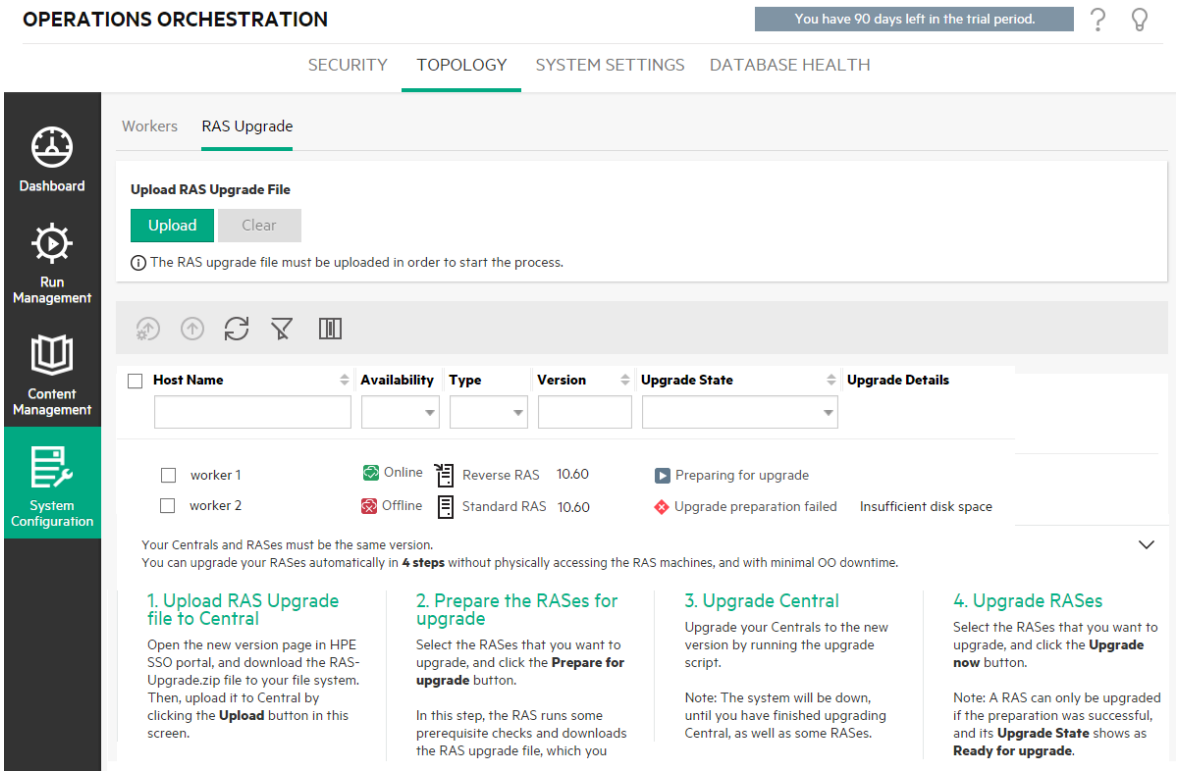

GUI item	Description
Upload	Opens the Upload File dialog box, so that you can select the relevant upload zip file to upload to Central.
Clear	Click this button to delete the RAS upgrade files from the Central database.
Prepare for Upgrade 	Click this button to get each selected RAS to perform the prerequisite checks and download the upgrader package from Central.
Upgrade Now 	Click this button to trigger the auto upgrade in the selected RASes.
Clear Filters 	Click this button to clear the filters and display all the RASes.

Refresh 	The RAS table is automatically refreshed every three seconds. However, you can also click this button to refresh the display.
Select Columns 	Click this button to adjust the display of columns.
Host Name	Displays the name of the host where the worker for this RAS is located.
Availability	Displays whether the RAS is online or offline.
Type	Displays whether the RAS is a standard RAS or reverse RAS. For more information about RAS types, see "Setting Up Topology – Workers and RASes" on page 81.
Version	Displays the version of HPE OO that the RAS is currently running.
ID	Displays the UUID (unique identifier) of the RAS.
Path	Displays the path where the worker for this RAS is located.
Upgrade State	Displays the status of the upgrade process for the RAS. • Preparing for upgrade • Upgrade preparation failed • Ready for upgrade • Upgrading • Upgrade failed • Upgrade succeeded
Upgrade Details	If there is a problem with the upgrade, this column displays the details of the problem. For example, if there was insufficient space for the upgrade file to be downloaded to the RAS.

Setting Up the System Settings - General Settings

In the general system settings (**System Configuration** > **System Settings** > **General**), you can set defaults and customizations, and configure a forward proxy, reverse proxy, or load balancer.

Customize elements of the HPE OO user interface

You can add extra text to the header title in the top banner in Central, in addition to the default text of “Operations Orchestration”.

By adding text, you can make Central look different in different environments, so that users will be able to identify which environment they are in. For example:

You can add a description of the ROI counter, in the Dashboard, in order to help users understand the meaning of the ROI values. For example, by entering the words "Total ROI in dollars" or "Total ROI in work days" you can show users that ROI is calculated in dollars or in work days.

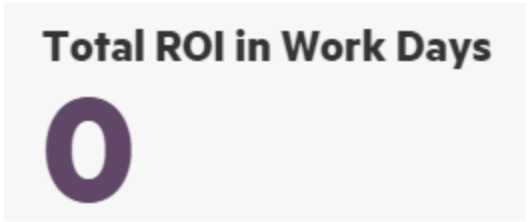

> **Note:** The text that you enter as the ROI description overwrites the default title of "Total ROI". Note also that this text does not get translated if you change the locale of the browser.

Set default timeout settings for runs

You can configure the default timeout settings for the system, so that if a flow run exceeds the specified duration, it will be canceled and will show the status **Canceled-Timeout Expired**.

> **Note:** The system timeout setting is applied to all flows in the system. However, if a flow is given a custom timeout setting in the Flow Library, or set to have none, this will overwrite the default. For more information, see "Managing the Flow Library" on page 149.

Configure the external URL

If you are using a load balancer, reverse proxy, or DNS load balancer, you need to tell HPE OO where the relevant external URL is located.

For example, **https://my.server.com:443/oo**.

If the URL of the load balancer host is not configured, the Run Tree and the Run Log will refer to the IP/host from the HTTP request.

Compression should not be set on the load balancer configuration. This can affect the connection between Central and workers and can lead to problems in new worker registration or existing worker functionality.

Configure a forward proxy

If you are using a forward proxy, you need to tell HPE OO its IP address, port number, username and password.

What do you want to do?

Add extra text to the header title

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **System Settings**.
3. In the **Extra text in title** box, enter the extra text to be displayed in the header title after “Operations Orchestration”.

   The maximum number of characters that can be entered is 30 for Latin letters and 19 for double-byte letters.
4. Click **Save**.

Add a description of the ROI counter

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **System Settings** > **General**.
3. In the **Override ROI description** box, enter a description of the value that is used to calculate ROI. For example, "Total ROI in US dollars".
4. Click **Save**.

Enable/disable the rerun functionality

1. Select the **System Configuration > System Settings** tab.
2. In the Run Settings area, select the **Enable the rerun option for failed or canceled runs** check box.

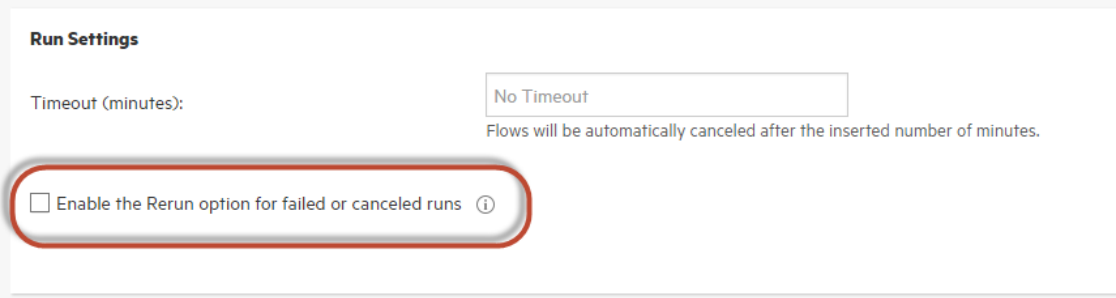

3. Reselect the check box to disable the rerun functionality. The rerun icon will no longer appear next to steps that can be rerun.

Configure the default timeout settings

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **System Settings** > **General**.
3. In the **Timeout (minutes)** box, enter a number larger than 0, to set the number of minutes for the default timeout setting. A flow run will be canceled if it exceeds this duration.

> **Note:** This value must be an integer.

4. Click **Save**.

Configure a external URL for a load balancer, reverse proxy, or DNS load balancer

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **System Settings** > **General**.
3. In the **URL** box, enter the URL of the load balancer, reverse proxy, or DNS load balancer.

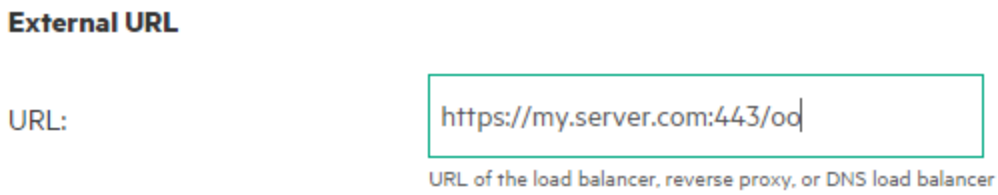

4. Click **Save**.

Configure a forward proxy

1. Click the **System Configuration** button to display the System Configuration workspace.
2. Select **System Settings** > **General**.
3. Enter the IP address, host number, username, and password for the forward proxy.

4. Click **Save**.

Reference Material

System Settings> General tab

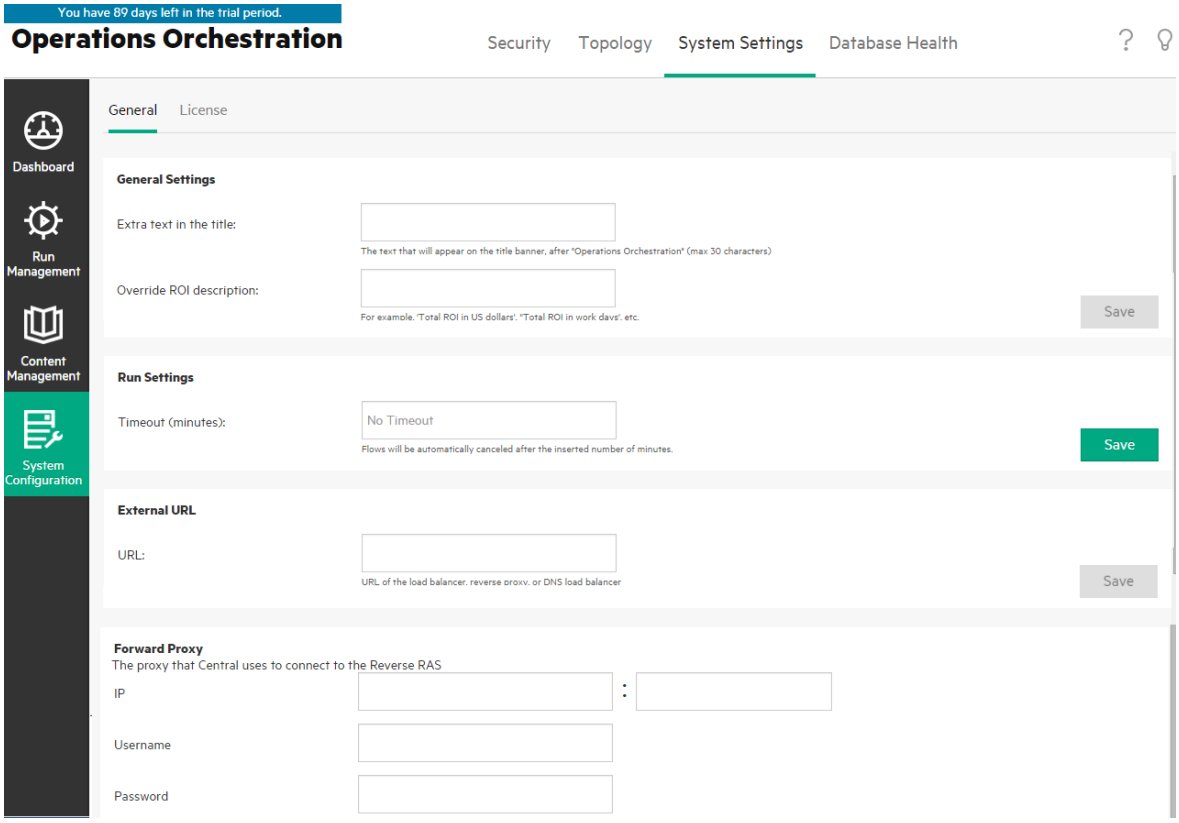

GUI item	Description
General Settings	
Extra text in the title	Enter the extra text to be displayed in the header title after “Operations Orchestration”. The maximum number of characters that can be entered is 30 for Latin letters and 19 for double-byte letters.
Override ROI description	Enter a description of the value that is used to calculate ROI. For example, "Total ROI in US dollars".
Run Settings	
Timeout (minutes)	Enter a number larger than 0, to set the number of minutes for the default timeout setting. A flow run will be canceled if it exceeds this duration.

	If nothing is entered, the system timeout is disabled. **Note:** This value must be an integer.
External URL	
URL	Enter the URL of the load balancer, reverse proxy, or DNS load balancer.
Forward Proxy	
IP	Enter the IP address and port number for the forward proxy.
Username	Enter the user name for the forward proxy.
Password	Enter the password for the forward proxy.

Setting Up the System Settings - Licensing

When you install or upgrade to HPE OO, this installs the Trial version license. You will need to install an Enterprise Edition license within 90 days.

HPE OO offers the following options for licensing:

- **HPE 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 HPE OO Enterprise Edition license to use the content packs. This the default license that comes with the HPE OO installer.
- **HPE OO Enterprise Edition**: A full license, unlimited in time, which includes all out-of-the-box content packs. This license unlocks the HPE OO Community Edition and HPE OO Enterprise Edition Trial licenses.
- **HPE OO Community Edition**: A partial license, which includes the Base content pack and the CE content packs, and enables 500 runs per month.

You can view the license information by clicking the **Info** button in the top right corner, 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.

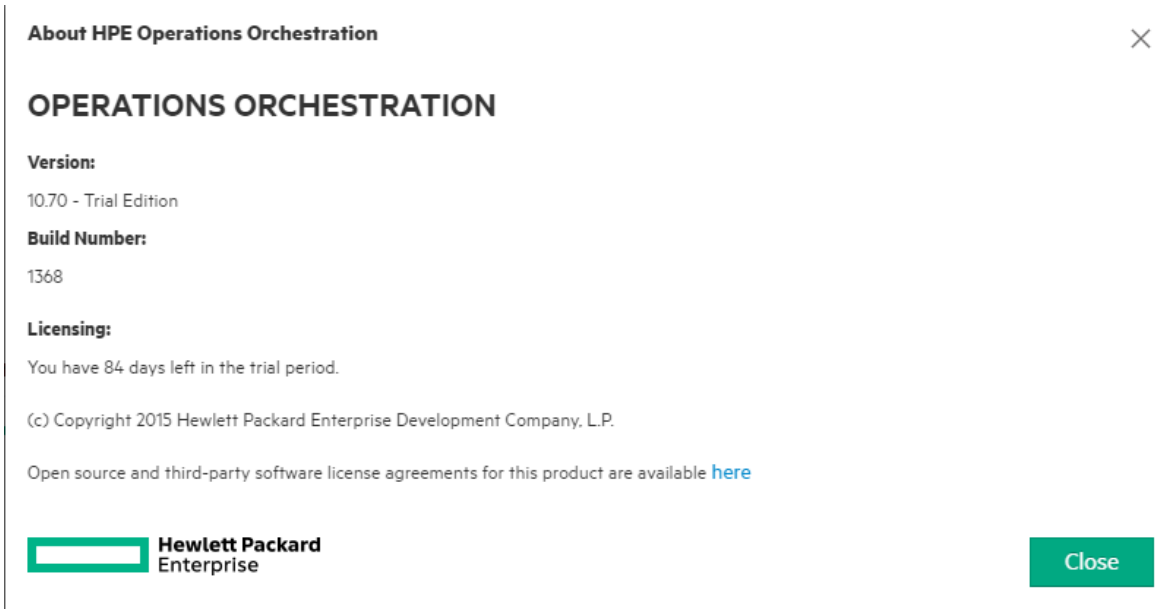

If you try to run flows after the license has expired, an error message appears, suggesting that you install an Enterprise Edition license.

You can also view details about the current license in the **License** tab under **System Configuration** > **System Settings**. From there, you can install a new license.

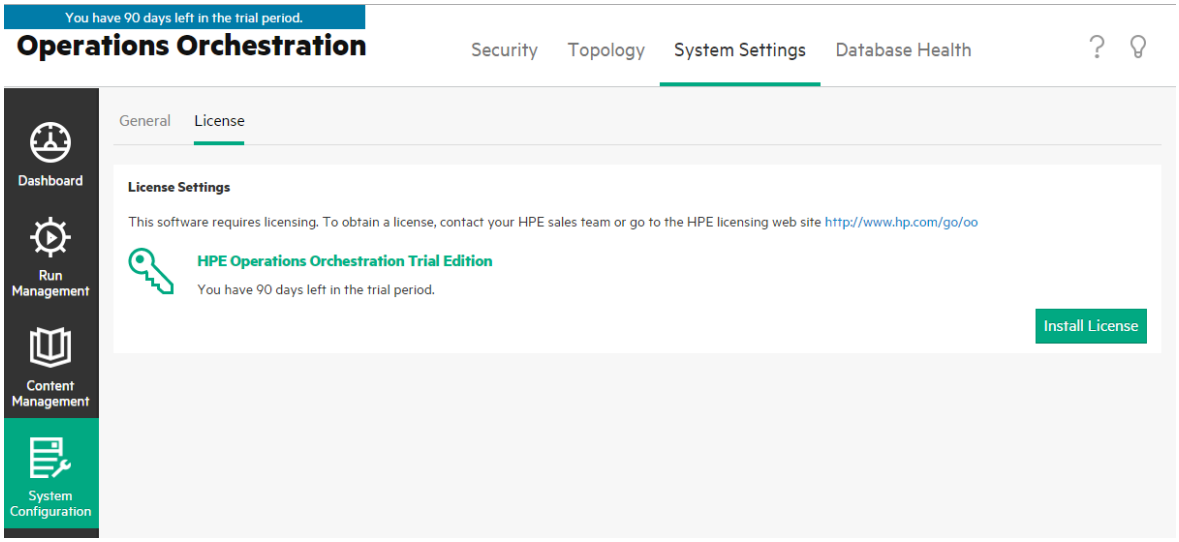

You can also see the remaining days/runs in the banner that appears at the top of all the Central windows.

You can click this banner from any location in Central, to open the **License** tab.

If you upgrade to a new version, the license is still valid.

Licensing a cluster

For a cluster, you need to request the license for one of the cluster IPs (and not on the load balancer IP) and install the license on that same cluster node.

This sets the license for the entire cluster even if that node is subsequently removed from the cluster.

What do you want to do?

Acquire a license

There are two ways to purchase a license:

- **Option 1**: Download the license from the HPE 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.

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

- **Option 2**: Issue a license using the HPE OO License Management system:
  - For a standalone Central installation, issue the license with the IP address of the Central server.
  - If you have CSA and HPE OO installed on the same machine, issue one license for that IP address.
  - If you have CSA and HPE OO installed on separate machines, issue two licenses for the two IP addresses.
  - For a cluster, choose one of the nodes and request a license for the IP address of this node. Then, install the license on that node

    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 HPE 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 HPE OO.

Install a license

1. Acquire the relevant license, as described above.

2. In Central, click the **System Configuration** button to display the System Configuration workspace.

   Important! If you are working with a cluster or if you have CSA and HPE OO on different machines, make sure that the open Central is at the IP address for which the license was issued.

3. Select **System Settings** > **License**.
4. Click the **Install License** button.
5. In the Install License dialog box, click **Browse** to navigate to the location of the license file, which was provided when you purchased the license.

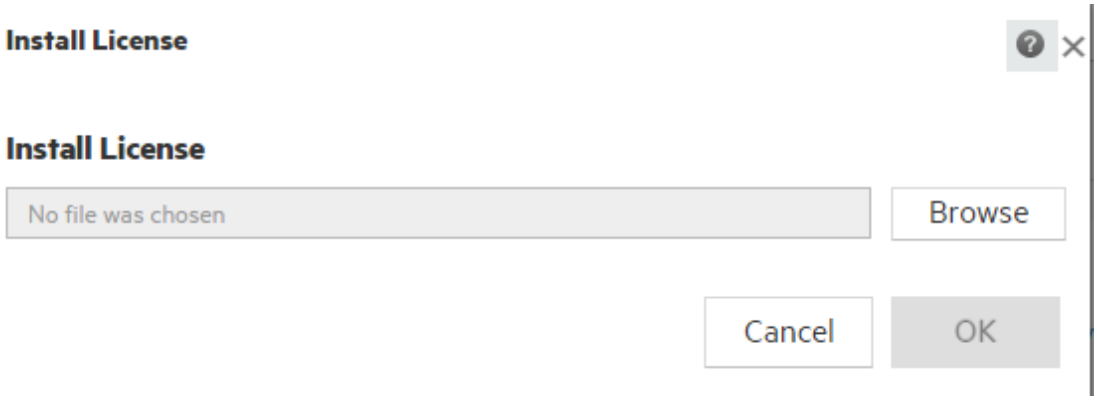

6. Click **OK**.

Note: You can also install a new license via the HPE OO Shell utility or via API. For more information, see the HPE OO *Shell (OOSH) User Guide* or the *HPE OO API Guide*.

Troubleshooting

In some cases, you may receive an error message saying that the license file is invalid.

If this error message appears, it could be because of the following:

- If you have **CSA and HPE OO installed on separate machines**, you issued two licenses for the two IPs.
  - Make sure that you selected the license for the HPE OO IP and not the CSA IP.

- For a **cluster**, you chose one of the nodes and issued a license for the IP address of this node.
  - Make sure that you are working from the Central at the IP address that was specified when you issued the license.
  - Make sure that you are not working from the Load Balancer IP.
- If you have **CSA and an HPE OO cluster**, you issued two licenses: one for CSA and one for a selected node.
  - Make sure that you are working from the Central at the IP address that was specified when you issued the license.
  - Make sure that you are not working from the Load Balancer IP.
  - Make sure that you selected the licence for the HPE OO IP and not the CSA IP.

Reference Material

License tab

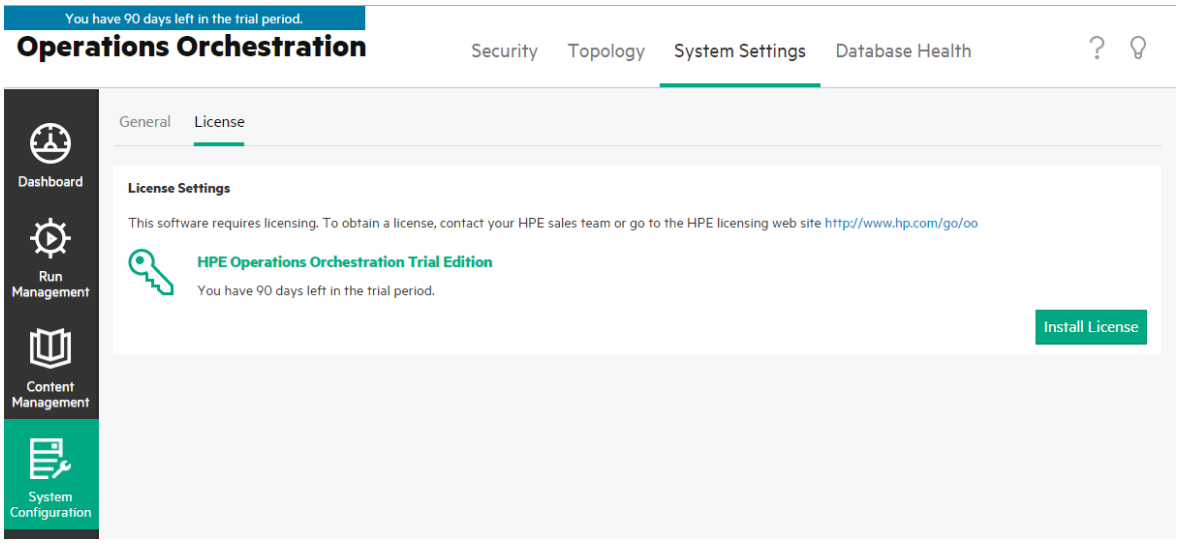

GUI item	Description
Install License	Click to display the Install License dialog box.

Install License dialog box

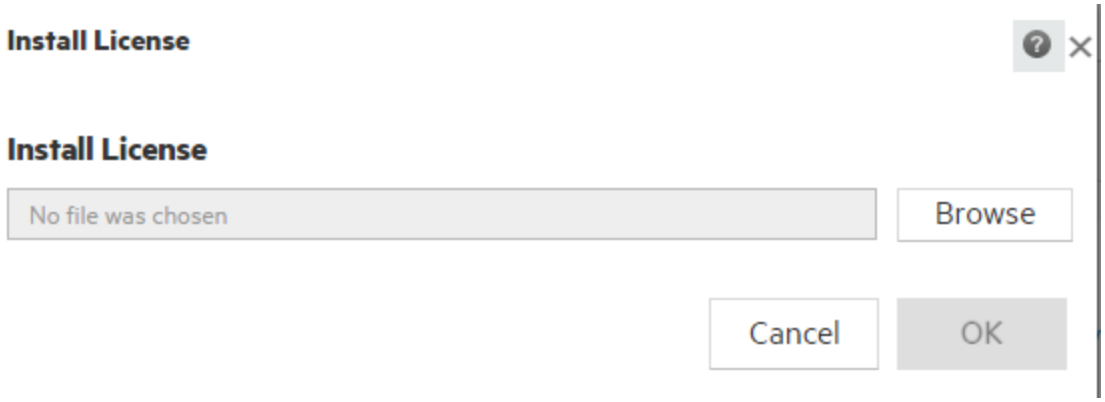

GUI item	Description
Browse	Click to navigate to the location of the license file, which was provided when you purchased the license.

Monitoring and Controlling Database Size

It is possible to monitor and control the size of your database, in order to enhance performance.

Under the **Database Health** tab in the System Configuration workspace, you can:

- Monitor the size of your database over different periods of time, and see how data is divided between execution, debugger and auditing usage.
- See information about how to schedule purging flows, in order to reduce the database size.
- Set the default level of information that is persisted in the Run Log. Reducing this level can help reduce the database size.

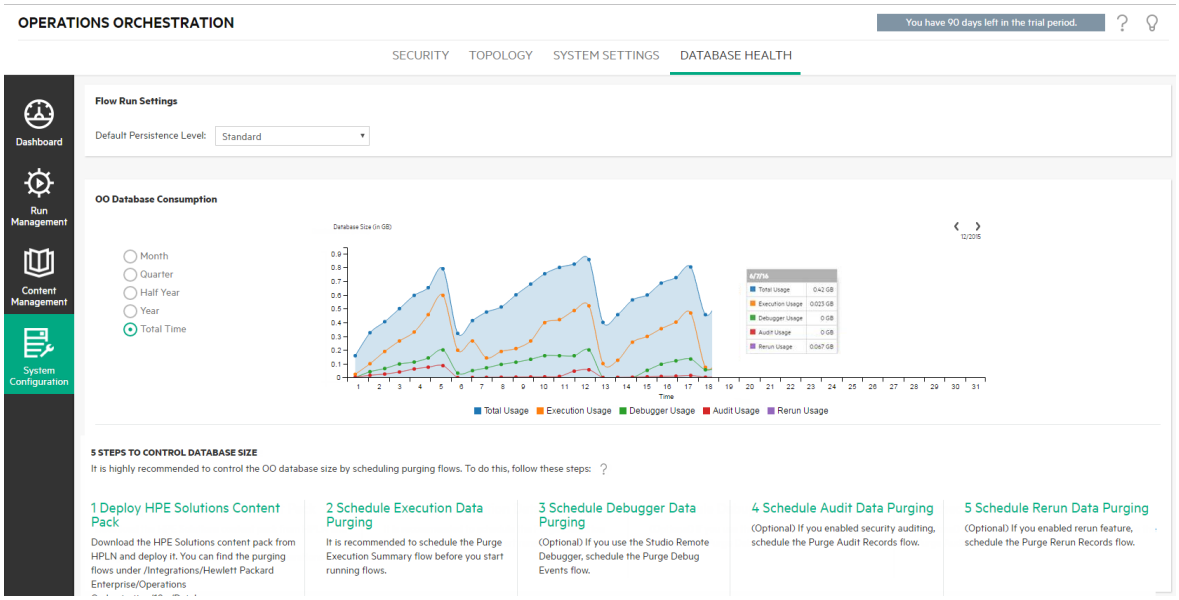

Monitor the Database Size

In the **Database Size** section, you can monitor the size of your database, comparing the current size to the size that it was in the past. You can adjust the graph to display the database size over the past month, quarter, half year, year, or total time. You can also see how the data is divided between execution, debugger and auditing usage.

> **Note:** The graph is updated approximately every 24 hours from the time Central was first up. For example, if you run a purging flow to reduce the database size, the updated size will be displayed

on the following day.

The **OO Database Consumption** section displays information about how to schedule purging flows, in order to reduce the size of the database. It is highly recommended to follow the described steps to deploy the **HPE Solutions** content pack and run the purging flows.

A Help  button is available, which you can click to see more information about scheduling purging flows.

Important! You should schedule the **Purge Execution Summary** flow before you start to run flows in Central. This will help reduce the size of the database and thus enhance the performance of HPE OO.

Note: The Step Persistence type selected for a step in a flow (Concise or Detailed) also affects the database size.

See "**Step Inspector > Advanced tab**" in the *HPE OO Studio Authoring Guide* for details on how to select the Step Persistence type and "Run Info > Step Details tab (Detailed Step Persistence)" on page 243 and "Run Info > Step Details tab (Concise Step Persistence)" on page 241 for details on how this effects the flow run.

Default Persistence Level for the Run Log

In HPE OO 10.x, a very detailed run history is saved to the database. This makes it easy to troubleshoot, as all the information is available in the Run Log. However, you may want to reduce the level of information that is persisted. For example, if your database size has increased to the limit or if some of the persisted data is not relevant for you.

By selecting the default persistence level, you can control the level of detail that is saved to the Run Log. You can choose between:

- **Extended** - large input/output values are not truncated when persisted.
- **Standard** - when persisted, large input/output values are truncated when they exceed the truncation threshold (in English, 4,000 bytes) in UTF-8 encoding.

  Note: The **Extended** level is similar to the behavior of previous 10.x versions. Therefore, when you upgrade from a previous version, the default persistence level is **Extended**, while in a new installation, the default is **Standard**.

  The top-level inputs and outputs of the launched flow are always logged with their full values (even at **Standard** level).

The default persistence level is applied to all flows that are run. You can override this default for individual flows in the flow library, or when triggering or scheduling a flow run.

> **Note:** All runs started from the Remote Debugger have the **Extended** persistence level.

How Values are Truncated at Standard Level

At **Standard** level, large input/output values are truncated when they exceed the truncation threshold in UTF-8 encoding. For English letters and symbols, this is 4000 bytes, but for other collations it may vary from 1000 to 4000 bytes (3000, in most cases).

This translates to:

- 4,000 ASCII characters
- 2,000 Latin-based characters (non-ASCII)
- 1,333 common CJK characters (Chinese, Japanese and Korean)
- 1,000 rare CJK characters and symbols

If the value combines different types of characters, the truncation threshold will be somewhere in between the above numbers. For example, Latin-based languages use a lot of ASCII characters, so the truncation point for such texts can be anywhere between 2,000 and 4,000 characters (typically over 3,000).

When the value is truncated, a “(truncated)” suffix is added to the stored value, so there are 3,985 bytes of actual content (instead of 4,000).

What do you want to do?

Set the persistence level for the Run Log

1. Click the **System Configuration** button to display the System Configuration workspace, and select the **Database Health** tab.
2. Select one of the options from the **Default persistence level** list:
   - **Extended** - large input/output values are not truncated when persisted
   - **Standard** - when persisted, large input/output values are truncated when they exceed the truncation threshold (in English, 4,000 bytes) in UTF-8 encoding.

This default level will be applied to all flows that are run. You can override this default for individual flows or folders in the flow library, or when triggering or scheduling a flow run.

Overwrite the persistence level for the Run Log

It is possible to overwrite the default persistence level that is set on the **Database Health** tab, in the following places:

- **Flow Library** - You can set the persistence level for individual flows in the Flow Library. For more information, see "Managing the Flow Library" on page 149.
- **Flow Launcher** - You can set the persistence level for a specific flow, when triggering a run. By default, the persistence level is taken from the system configuration (first) and the Flow Library (second). For more information, see "Running a Flow" on page 181
- **Scheduler** - You can set the persistence level for the specific scheduler configuration. By default, the persistence level is taken from the system configuration (first) and the Flow Library (second). For more information, see "Scheduling Flow Runs" on page 196.

Monitor the size of your database

1. Click the **System Configuration** button to display the System Configuration workspace, and select the **Database Health** tab.
2. Under **OO Database Consumption**, select one of the options to display the database size (in MB) over the past month, quarter, half year, year, or total time.
3. If required, click the arrows to display areas of the graph that are hidden.
4. Click the Legend blocks under the graph to filter the display, in order to hide/show the execution, debugger, and auditing and rerun data.

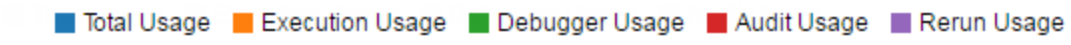

Control the size of your database

It is strongly recommended to schedule the purging flows, in order to control the size of the HPE OO database and enhance the performance of HPE OO.

Important! Make sure to schedule the **Purge Execution Summary** flow *before you start* to run flows in Central!

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

1. Download the **HPE Solutions** content pack from HPLN.
2. Deploy this content pack to Central.
   a. Open the **Content Packs** tab in Content Management and click the **Deploy New Content**  button.

b. In the Deploy New Content dialog box, click the **Add** + button.

c. Browse to and select the **HPE Solutions** content pack, and then click **Open**.

For more information, see "Deploying and Managing Content Packs" on page 127.

3. Open the **Flow Library** tab in Content Management and locate the purging flows, under **Integrations/Hewlett-Packard Enterprise/Operations Orchestration/10.x/Database**.
4. Open the **Scheduler** tab in **Run Management.**
5. Create schedules to run the following flows on a regular basis:
   - **Purge Execution Summary** - recommended for all users
   - **Purge Debug Events** - if you use the Studio Remote Debugger
   - **Purge Audit Records** - if you have enabled security auditing
   - **Purge Rerun Info** - if you have defined rerun points

   For more information, see "Scheduling Flow Runs" on page 196.

Reference Material

Database Health tab

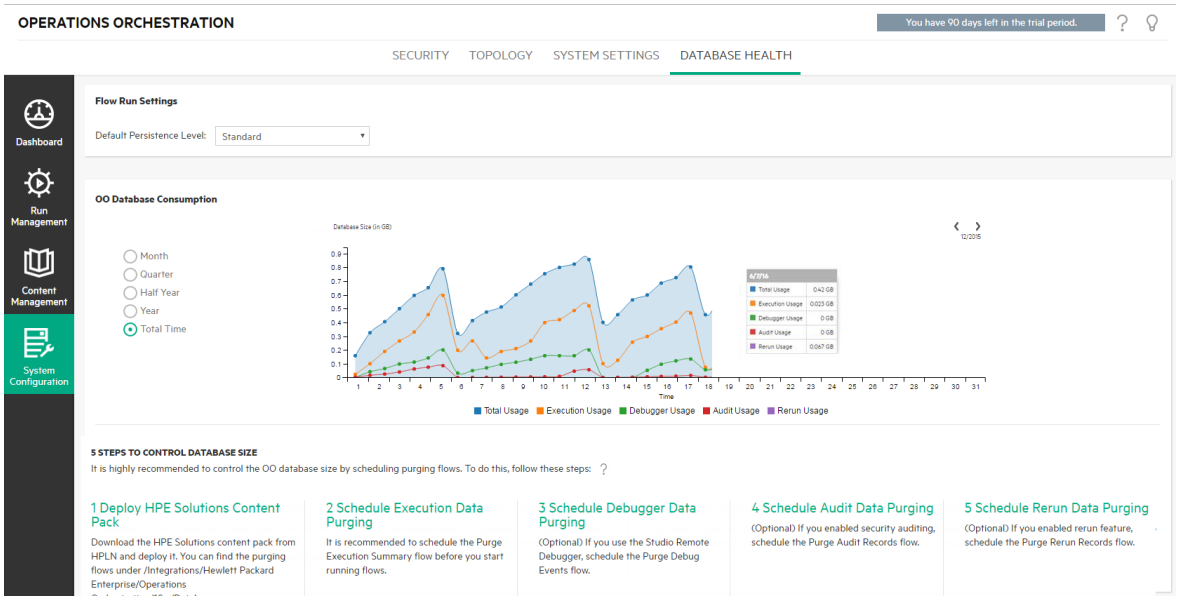

GUI item	Description

Run Log Persistence Level	Select the default level of detail that will be saved to the Run Log when a flow is run: • **Standard** - large input/output values are truncated at approximately 4,000 bytes when persisted • **Extended** - large input/output values are not truncated when persisted > **Note:** The **Extended** level is similar to the behavior of previous 10.x versions. > The top-level inputs and outputs of the launched flow are always logged with their full values (even at **Standard** level).
OO Database Consumption	See a graph showing the database size in MB, over a period of time, ending with the current date. Select the options to change the displayed time period. If required, click the arrows to display areas of the graph that are hidden.
Legend blocks	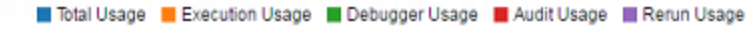 Click the Legend blocks under the graph to filter the display, in order to hide/show the execution, debugger, and auditing data.
5 Steps to Control Database Size	View the steps describing how to deploy the **HPE Solutions** content pack and run the purging flows.

Controlling Database Size by Purging

It is highly recommended to schedule the purging flows, in order to control the size of the HPE OO database and enhance the performance of HPE OO.

Important! Make sure to schedule the **Purge Execution Summary** flow *before you start* to run flows in Central!

1. Download the **HPE Solutions** content pack from HPLN.
2. Deploy this content pack to Central.
   a. Open the **Content Packs** tab in Content Management and click the **Deploy New Content** button.
   b. In the Deploy New Content dialog box, click the **Add** button.
   c. Browse to and select the **HPE Solutions** content pack, and then click **Open**.

   For more information about deploying content packs, see "Deploying and Managing Content Packs" on page 127.
3. Open the **Flow Library** tab in Content Management and locate the purging flows, under **Integrations**/**Hewlett Packard Enterprise**/**Operations Orchestration**/**10.x**/**Database**.
4. Open the **Scheduler** tab in the Run Management workspace.
5. Create schedules to run the following flows on a regular basis:
   - **Purge Execution Summary** - recommended for all users
   - **Purge Debug Events** - if you use the Studio Remote Debugger
   - **Purge Audit Records** - if you have enabled security auditing
   - **Purge Rerun Info** - if you have defined rerun points

   For more information about scheduling, see "Scheduling Flow Runs" on page 196.

For more information about using the purging flows, see the HPE OO 10 Tutorials section, available on HPLN here.

Setting the Browser Language

Central language support is according to the browser language. If the language is not supported, Central is displayed in English.

Set up your browser for foreign language support.

> **Note:** Make sure that the added language is the first in the list.

Change the language preferences on Internet Explorer

1. Go to **Tools** > **Internet Options** > **General** (tab) > **Languages** > **Language** preference.
2. Select the required language. Make sure the primary language is the first in the list.

Change the language preferences on Firefox

1. Go to **Tools** > **Options**, **Content** (tab) > **Languages** > **Choose** (button).
2. Select the required language.

Change the language preferences on Google Chrome

1. Click the **Customize and control Google Chrome** button, and then select **Settings** > **Show advanced settings**.
2. In the **Languages** section, click **Language and input settings**.
3. Click **Add**, and select the required language.

   You can add languages to this from a predefined set, and Chrome will send all the languages you choose with the Accept-Language header in the order in which you arrange them.

Change the language preferences on Safari

In Windows, you can't change the language sent by Safari. The language sent with the Accept-Language header is derived from your system settings.

> **Note:** It is also possible to change the MS SQL collation language or the **central-wrapper.conf** language for content. For more information, see the *HPE OO Installation, Upgrade, and Configuration Guide*.

Working with CloudSlang Content Flows

 It is also possible to deploy and run content that was authored using CloudSlang 1.0.

CloudSlang

CloudSlang is a flow-based orchestration tool for managing deployed applications.

The CloudSlang project is composed of three main parts: the CloudSlang Orchestration Engine, the CloudSlang language and the ready-made CloudSlang content.

The engine is packaged as a lightweight Java .jar file and can be embedded into existing Java projects.

The language is a YAML-based DSL for writing workflows. The workflows can be run by an embedded instance of the engine or the stand-alone CloudSlang CLI. The main types of CloudSlang content are operations, flows, and decisions. An operation contains an action, which can be written in Python or Java. Operations perform the “work” part of the workflow. A flow contains steps, which stitch together the actions performed by operations, navigating and passing data from one to the other based on operation results and outputs. Flows perform the “flow” part of the workflow. A decision provides the capability to select different outcomes.

The supported file extensions are: .sl, .sl.yaml and .sl.yml.

> **Note:** For more information about CloudSlang, see:
>
> http://www.cloudslang.io/#/docs
>
> https://github.com/cloudslang/cloud-slang
>
> https://github.com/cloudslang/score

You can run CloudSlang flows in all the different HPE OO topologies—single node, Central and RAS, and clustering.

Version HPE OO <10.70 can run CloudSlang <1.0 and HPE OO 10.70 can run CloudSlang 1.0. See the HPE OO *Release Notes* for the supported version of CloudSlang.

Content Lifecycle when working with CloudSlang

CloudSlang content has the same life cycle as OO content:

- Deploy a CloudSlang content pack. See "Deploying and Managing Content Packs" on page 127.
- Manage Flows: View flows in the flow library and set flow permissions. See "Managing the Flow Library" on page 149
- View and edit System Properties. For CloudSlang, System Properties can also be configured as Sensitive.See "Setting Up Configuration Items for a Content Pack" on page 164.
- Run and monitor flows: Run flows, schedule runs, track and manage flow runs. See "Running and Monitoring Flows" on page 179.
- Troubleshooting Flows Runs: Drill down into the flow run. See "Testing and Troubleshooting a Flow Run" on page 229.

Features supported by CloudSlang and/or OO Flows

This section describes where CloudSlang content and OO Flows content behavior is different:

- **Content Permissions.** Permissions to run a CloudSlang flow is are computed when the flow is launched or resumed. If you don't have the correct content permissions for all dependencies, you will not be able to launch or resume the flow. See "Setting Permissions for Content" on page 150
- **Flow graphical representation.** A flow graphical representation is not available: If you open the Flow Graph tab for a flow that was written in CloudSlang, a **not supported** message is shown. See "Viewing a Flow Graph" on page 159.
- **Sensitive Data for System Properties.** In CloudSlang, OO flows, also CloudSlang flows can include Sensitive Inputs and Outputs, and they will be hidden in Central.
- **Decisions.** CloudSlang step type called **Decision**, works along with existing subflows and operations. When running a flow that includes a decision, the result of a decision looks like a Custom result.
- **Content Distribution.** All required resources are resolved and downloaded from Central to the RAS at the start of the operation execution using maven. For Python based content, the administrator must add manually, in the python library in each RAS, all the dependencies needed to run the flow and are not included in HPE OO out-of-the-box python lib. This is done by adding them to the HPE OO class path in ALL HPE OO nodes in your environment (Central or RAS):
  - For Central: **<Installation_path>\central\lib\pythonlib**
  - For RAS: **<Installation_path>\RAS\lib\pythonlib**
- **CloudSlang content IDs** use the format **<namepace>.<name>**.

Feature	Supported in CloudSlang	Supported in OO Flows
Content Packs. Deploy, rollback, delete, view details in the Content Management workspace. See "Deploying and Managing Content Packs" on page 127.	V	V
Distribution Capability.	V Java automatically V Python manually	V
Content Permission. See ."Managing the Flow Library" on page 149	V	V
Run Flows. See "Running a Flow" on page 181	V	V
Schedule a flow. See "Scheduling Flow Runs" on page 196.	V	V
Monitor Flows. See "Adjusting the Display of Flow Runs" on page 223	V	V
Manage Flow Runs. See"Tracking and Managing Flow Runs" on page 215.	V	V
Decisions.	V	X
Sensitive Data in Input and Output Fields	V	V
Sensitive Data for System Properties	V	X
View the Flow Graph for a flow. See "Viewing a Flow Graph" on page 159.	X	V
Worker Group. See Worker Groups and Group Aliases in the *HPE OO Concepts Guide*.	X	V

Promoting Content Packs

The Op Admin is usually the one responsible for promoting content packs.

Promoting a Content Pack – Overview

What is promotion?

Promotion is the process of deploying content packs across several environments, in order to ensure that results are predictable and without (unknown) risks. This is in particular important for maintaining a stable production environment that contains content that has been tested and validated prior to deployment.

For example, deploying the content pack on the following Central environments: Development, QA, Staging, Production.

The final aim of promotion is to deploy a new content pack to the Central Production server, to make the flows available to users.

Is promotion the same as deployment?

Deployment is part of promotion. However, the promotion process also includes other tasks, such as:

- Configuring the content pack: configuring worker group aliases, mapping system accounts, and so on
- Testing and troubleshooting the flows in the content pack

What is a content pack?

A content pack is a file containing operations, flows, actions, and configuration items. When flow authors complete a project, they package it into a content pack for promotion on the Central server.

- CloudSlang operations use actions that are Java-based or python-based.
- Native OO operations use actions that are Java-based or .Net based.

For more information about content packs, see the *HPE OO Concepts Guide*.

When do you promote a content pack?

You promote a content pack when:

- There is a need for new content in the Production environment.
- You have fixed a bug in a flow and need to upload the new version of the flow to the Production environment.
- You need to add new functionality to an existing flow in the Production environment.
- You are upgrading to a new version of an out-of-the-box content pack.

Who does the promotion?

The promotion process in each of the environments is slightly different and may be done by a different persona. However, in most cases, it is the Op Admin who is responsible for promotion.

The person doing the promotion must be assigned a role with **View Content Packs** and **Manage Content Packs** permissions.

For more information, see the *Best practices for HPE OO content promotion* document, available on HPLN.

Promotion steps

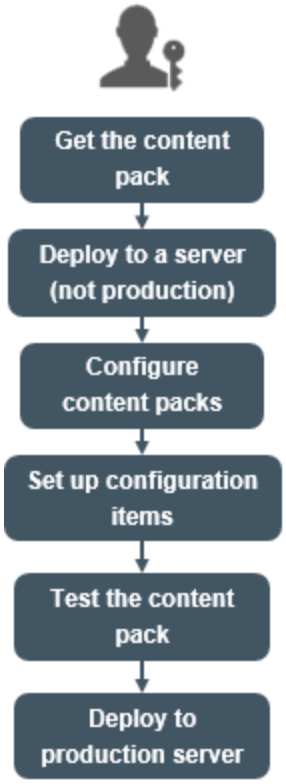

> **Note:** This is just a high level look at the Central workflow, and there are many options that are not described here. For more detailed information about any of the steps, use the links to learn about the options in-depth.

Step 1: Get the Content Pack From the Flow Author

1. The Flow Author creates a content pack in Studio. See "Exporting a Content Pack" in the *Studio Authoring Guide*.
2. The Flow Author saves the content pack to the Artifact Repository.
3. The Flow Author shares the content pack with the Op Admin/Promoter.

Step 2: Deploy to a non-Production Server

1.  The Op Admin receives the content pack from the Flow Author.
2.  The Op Admin deploys the content pack on the a non-Production server, for example, the Staging server. See "Deploying and Managing Content Packs" on the next page.

Step 3: Configure the Content in the Content Pack

In this step, the Op Admin adjusts the content pack to the environment by configuring the content in it.

1.  The Op Admin sets the content permissions on the flows. See "Managing the Flow Library" on page 149.
2.  The Op Admin sets the persistence level and run timeout for the flows. See "Managing the Flow Library" on page 149.

Step 4: Set up the Configuration Items in the Content Pack

In this step, the Op Admin adjusts the content pack to the environment by setting up the configuration items in it.

1.  If the content pack includes system and system properties, the Op Admin assigns values to these in the content pack. See "Setting Up Configuration Items for a Content Pack" on page 164.
2.  The Op Admin maps the aliases of the worker groups to actual worker groups. See "Setting Up Configuration Items for a Content Pack" on page 164.

Step 5: Test and Troubleshoot the Flows in the Content Pack

1.  The Op Admin runs each flow from **Content Management** > **Flow Library** and checks whether it runs properly. See "Tracking and Managing Flow Runs" on page 215.
2.  If a flow run fails, the Op Admin can drill down into a flow to identify the problems. See "Testing and Troubleshooting a Flow Run" on page 229.

Step 6: Deploy to the Production Server

1.  If required, the Op Admin adjusts the configuration of the content pack, for the Production server. For example, it may be necessary to map the worker group aliases and system accounts

differently for this server.

2.  Finally, the Op Admin deploys the content pack to the Production server.

Deploying and Managing Content Packs

 In the **Content Packs** tab, you can deploy content packs to the Central environment. For example, Development, Testing, Production. Deployment is performed at least twice in the promotion process.

> **Note:** It is advised to deploy the base content pack and any other HPE OO content packs that are used in Studio before deploying your own content packs.

Viewing Content Pack Details

After content packs have been deployed, they are displayed on the **Content Packs** tab.

> **Note:** After an upgrade, the deployment time is not displayed for content packs that were deployed to the previous version.

The **Content** tab on the right side displays more details about the selected content pack. This information includes the version, publisher, description, digital signature details, the name of the user who deployed it, the Source Control Management tag, SCM revision (for SVN), and the entities inside the content pack. The contents are arranged in a tree, under **Library** and **Configuration** folders, reflecting the way that the content pack appeared in Studio. You can drill down into the folders and see exactly which items exist in the content pack and where they are stored. This includes flows, operations, and configuration items.

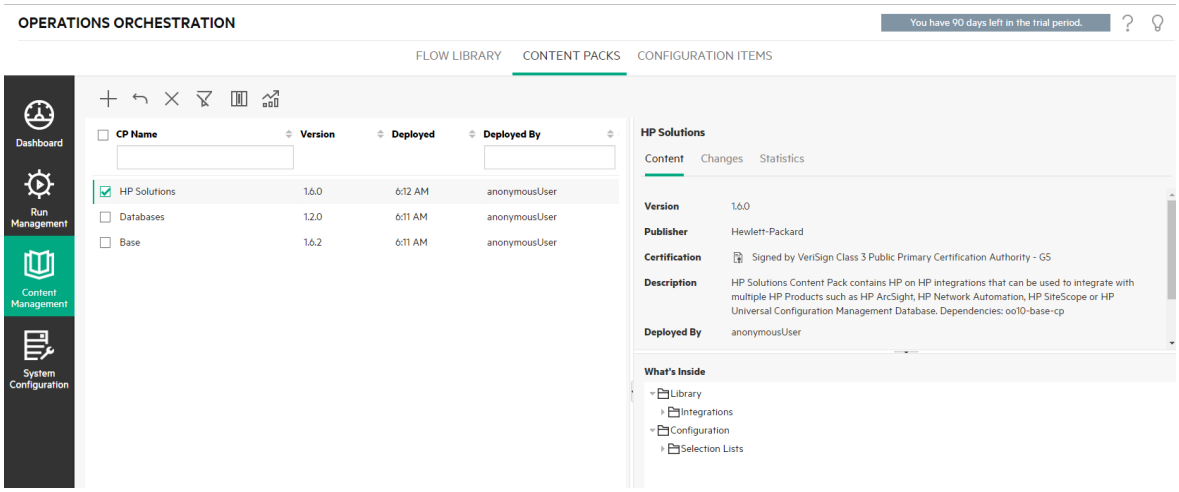

The **Changes** tab displays the changes that have been made to the selected content pack since the version that was previously deployed on Central. This information tells you which items have been added, deleted, updated, or moved.

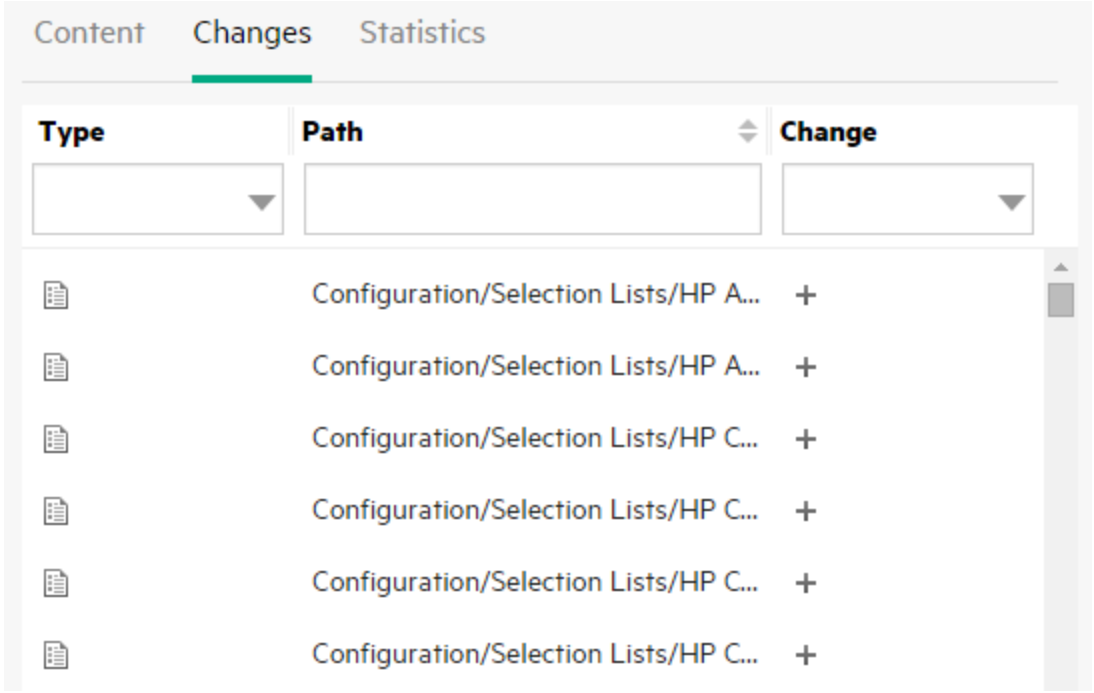

The **Statistics** tab displays statistical information about the selected content pack:

- See how many flows were run were using content from this deployed content pack for each month since first use, and compare this with the total number of all runs over the same period.
- See how this content pack usage fits into the percentage of all runs during the current month, previous month, quarter, half year, year, or in total.

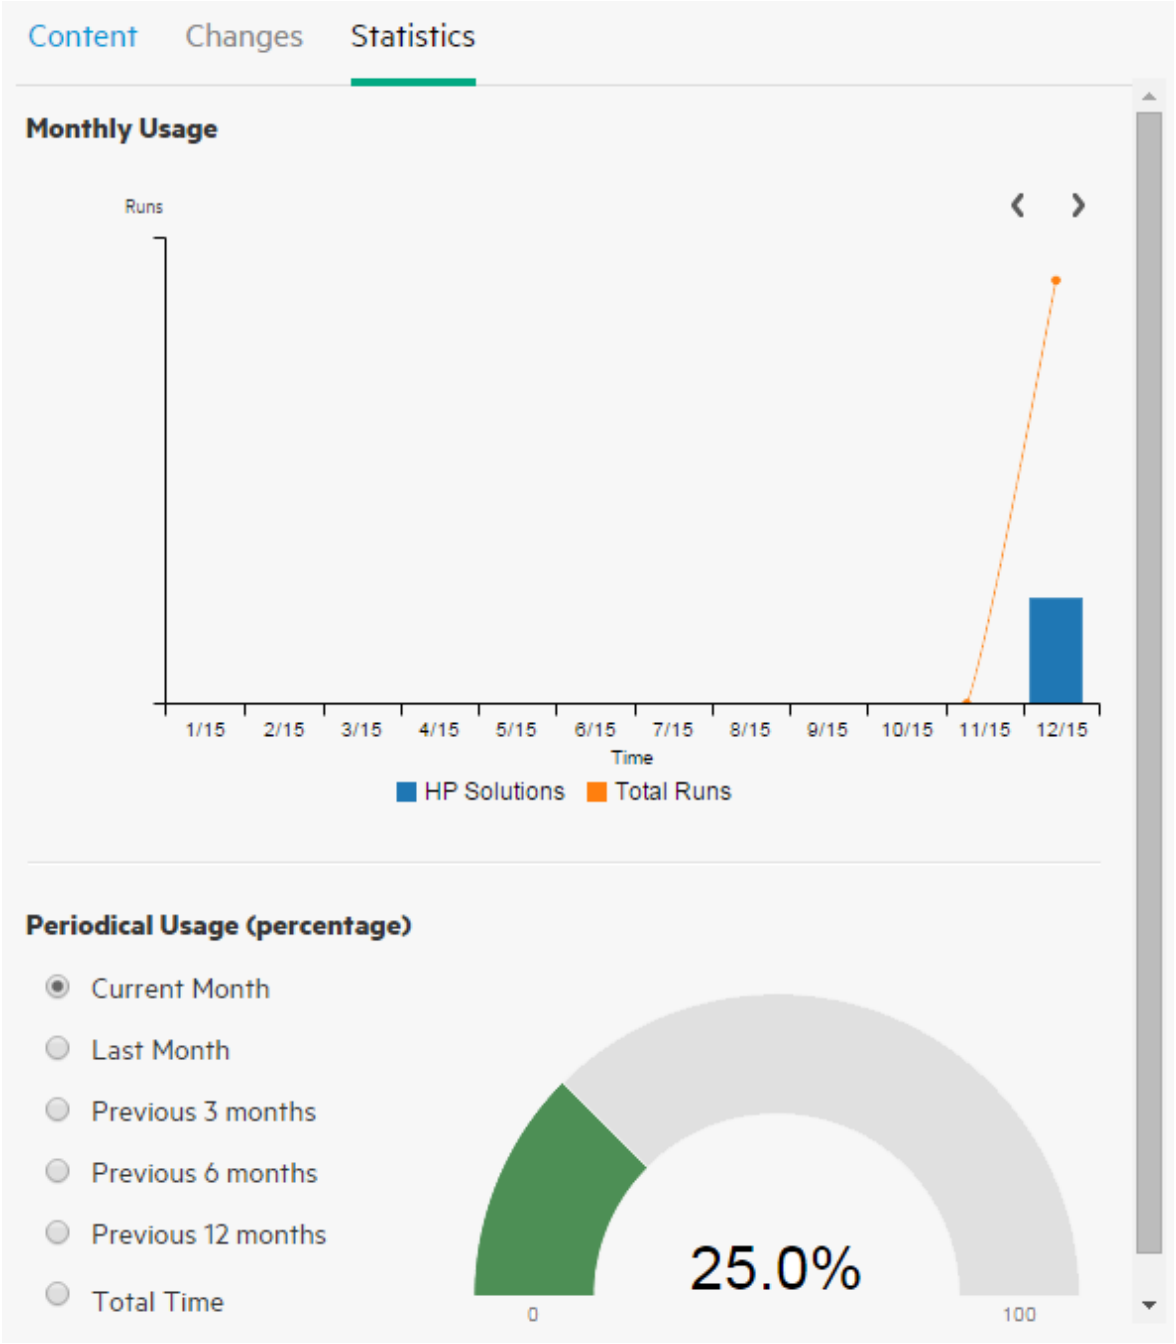

You can also click the Statistics button in the toolbar to display and compare statistical information about all deployed content packs.

Content Pack Versions

You can deploy a content pack to a server multiple times, as the content is updated by flow authors, and new versions of the content pack are created.

After a content pack has been deployed, it is possible to roll back to the last deployment version. This might be useful if you find that there are problems with the flows in a new deployment. For more information, see Roll back to a previous deployment.

> **Note:** When you roll back to the last deployment, you return the content to the state that it was in before your most recent deployment. If the most recent deployment involved multiple content packs, the rollback returns all of these content packs to the state they were in, previously.
>
> Note also that you can only roll back one version. You cannot roll back to earlier versions than the last one.

Deploying Multiple Content Packs

A single user can deploy multiple content packs at once, by using the **Add** + button in the Deploy New Content dialog box.

However, it is not possible for two users to deploy content packs to the same Central at the same time. If you try to deploy content to Central while another user is already deploying content, you will see an error message, indicating that another deployment is currently in progress.

Split Deployment

After a content pack has been deployed, it is possible to split the original content pack into two or more parts in Studio and deploy each of these separately. Note that if one of these smaller content packs has the same name as the one that was originally deployed, it will overwrite the original one on the Central server. If it has a different name, it will be deployed alongside it.

Content Pack Validation - Digital Signatures

A valid digital signature on a content pack verifies that the content pack was created by a known sender and that it was not altered in transit.

During the upload process, before deployment, Central checks that content packs with digital

signatures have not been altered or expired, and displays the details of the certificate.

> **Note:** HPE OO verifies the following common signature types: .DSA, .RSA, .EC, .SF

If a signed content pack has been altered, that content pack cannot be deployed. If the signature has expired, the content pack can be deployed only if you select a check box agreeing to ignore this expiry. For more information, see Deploy a content pack.

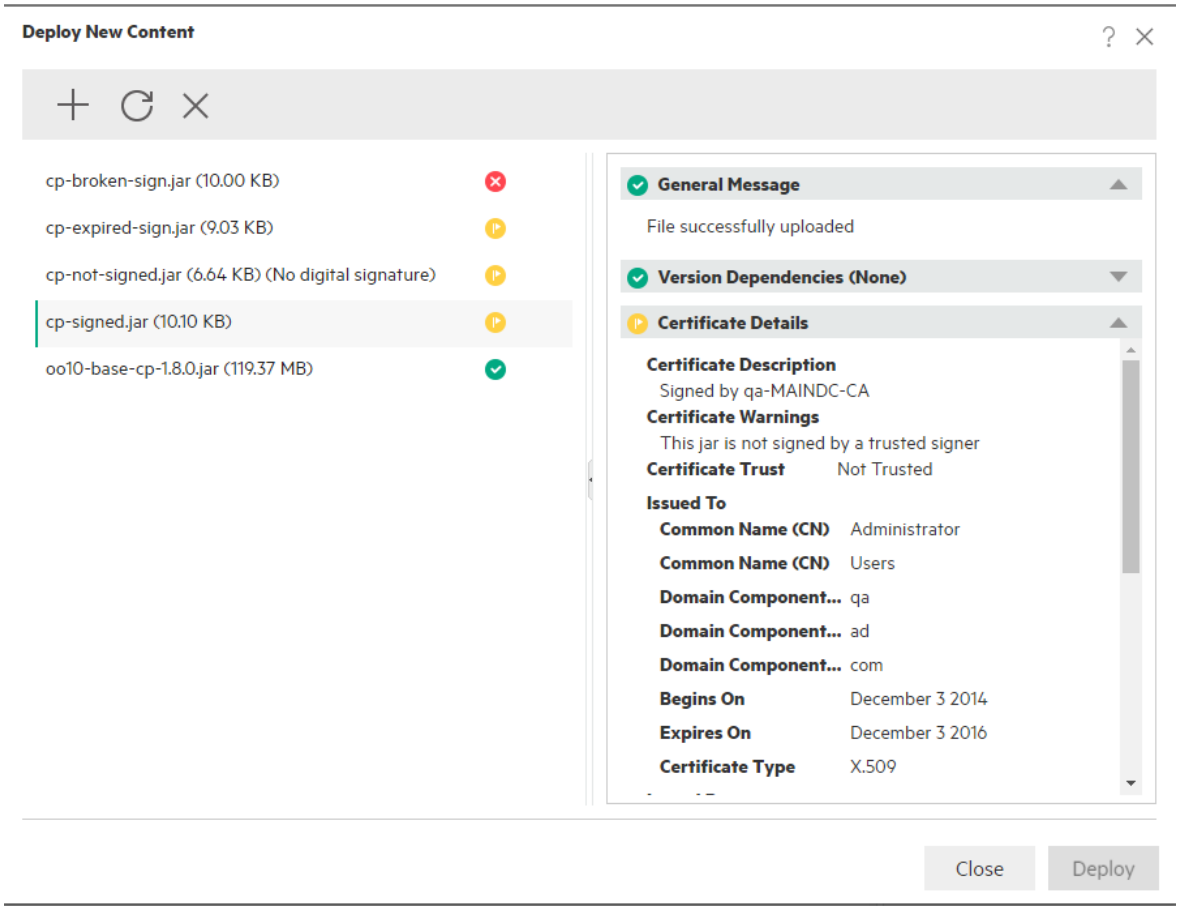

The scenarios for signatures are the following:

- The out-of-the-box content packs created by HPE OO are signed with the HPE private key.
- When you buy a content pack from a content pack vendor, the vendor uses an official certificate authority (CA) certificate.

  If the signing authority is not already in the Central TrustStore, you will need to import the CA certificate into it, in order to enable Central to verify the content pack's certificate. For details about how to add a CA to the Central TrustStore, see "Importing a CA Root Certificate to the Central TrustStore" in the *HPE OO Security and Hardening Guide*.

- When HPE OO authors create content packs in-house, they need a CA private key to be able to add a signature to a content pack. You can request this certificate from a company such as VeriSign, for example. It is also possible to use a self-signed signature, but this is not recommended, for security reasons.

  a. The promoter adds a digital signature to the content pack, using an external signing tool, such as Java JarSigner.

  b. The administrator adds the trusted root certificate authority (CA) into the HPE OO **client.truststore** file (under **central/var/security**). In a cluster, import the CA to the **client.truststore** file for all nodes.

     For details about how to add a CA to the Central TrustStore, see "Importing a CA Root Certificate to the Central TrustStore" in the *HPE OO Security and Hardening Guide*.

Security Note: It is recommended to verify the digital signature of a content pack manually using well known tools before importing it into Studio.

Content Pack Validation - Dependencies

During the upload process, before deployment, Central checks whether the content pack has any dependencies. If you are deploying a content pack that uses elements (for example, operations or properties) that are in a different content pack, the content pack that you are deploying will have a dependencies file that points to this other content pack.

However, what happens if any of these items were deleted from the other content pack, and the new version of this content pack has been deployed in your environment? If this occurs, the content pack that you are now trying to deploy will have missing dependencies, and this could cause the deployment to fail.

If there are missing dependencies, Central displays a warning icon in the Deploy New Content dialog box and displays the expected version of the content pack that contained these dependencies. This gives you an opportunity to add the expected version of this content pack to the deployment, so that the deployment of the content pack that depends on it will be successful.

Note, however, that you have the option to continue with the deployment without adding the dependencies.

For more information, see "Managing Content Packs and Dependencies in a Project" in the *HPE OO Studio Authoring Guide*.

Content Pack Validation - Duplications

During the deployment process, Central performs a complete path validation for all elements that have a unique path on Central (such as flow, group alias, domain term, system property, system account, and selection list).

If there are any duplications of these elements (either within the deployment unit, or within previously deployed entities), an error message is displayed, listing the duplications.

Content Pack Validation - Structural Integrity

During the deployment process, Central validates the content packs to make sure that the structural integrity of the deployed content is intact.

Checking the structural integrity includes checking that there are no missing references, there are no duplicate UUIDs, the flows in the content packs are correctly formed, and there are no duplicate flow paths or duplicate system property names.

For example:

Deployment After Making Changes to a Content Pack

Each entity (flow, operation, and configuration item) in deployed content packs must have a unique UUID. If you try to deploy a content pack with a duplicated entity (with the same UUID as an entity in another deployed content pack), the deployment will fail.

If you move an entity from one content pack to another, this may result in a duplication, which will cause deployment to fail. For example, you have two content packs, A and B. You deploy content pack A to Central, and then you move entity X from A to B. When you try to deploy content pack B, the deployment will fail because entity X appears in both A and B.

In such a situation, the correct workflow is to deploy ***both A and B*** together after moving the entity, and not to deploy B on its own. The new version of A, without entity X, overwrites the previous version on Central, so there is no duplication.

When you want to promote content from one Central environment to another, you will only need to deploy all of the latest versions of the content packs that were deployed in the original Central environment.

Best Practices for Content Pack Deployment After Upgrade

After an upgrade, there may be content deployed in previous versions that does not meet the requirements for uniqueness. If this occurs, there is an indication in the **Content Packs** tab, showing that the content pack may not be valid (some of its contents might have moved to a different content pack without a new version of the first content pack being re-deployed).

If you are upgrading a large repository from HPE OO 9.x, the recommended best practice is to:

1. Upgrade the repository to HPE OO 10.x.
2. Deploy it to Central.
3. After the repository has been successfully deployed, split it into multiple content packs in Studio.
4. Deploy those content packs together.
5. If you modify, remove, or move entities in these content packs, deploy all affected content packs again.

What do you want to do?

Deploy a content pack

When you deploy a content pack, HPE OO checks the content pack for a digital signature.

- If a content pack with a digital signature has been altered, this content pack cannot be deployed.
- If the signature has expired, you will need to select the **Ignore expiry of digital signature** check box in order to deploy the content pack.

- If there is no digital signature on the content pack, you can deploy it, but a warning is displayed, because there is no verification that the content pack was created by a known sender and that it was not altered in transit.
- If a content pack is signed, but the signature is not from a known and trusted authority, you can still deploy it.

1. Click the **Content Management** button to display the Content Management workspace.
2. Under the **Content Packs** tab, click the **Deploy New Content**  button.
3. In the Deploy New Content dialog box, click the **Add**  button.
4. Browse to and select one or more content packs, and then click **Open**.

   > **Tip:** You can also drag-and-drop a content pack from your file management system into the Deploy New Content dialog box.

5. During the upload process, OO verifies the selected content packs. Icons next to the content packs indicate if there are any issues with their certification.

   Select a content pack to display its certification details.

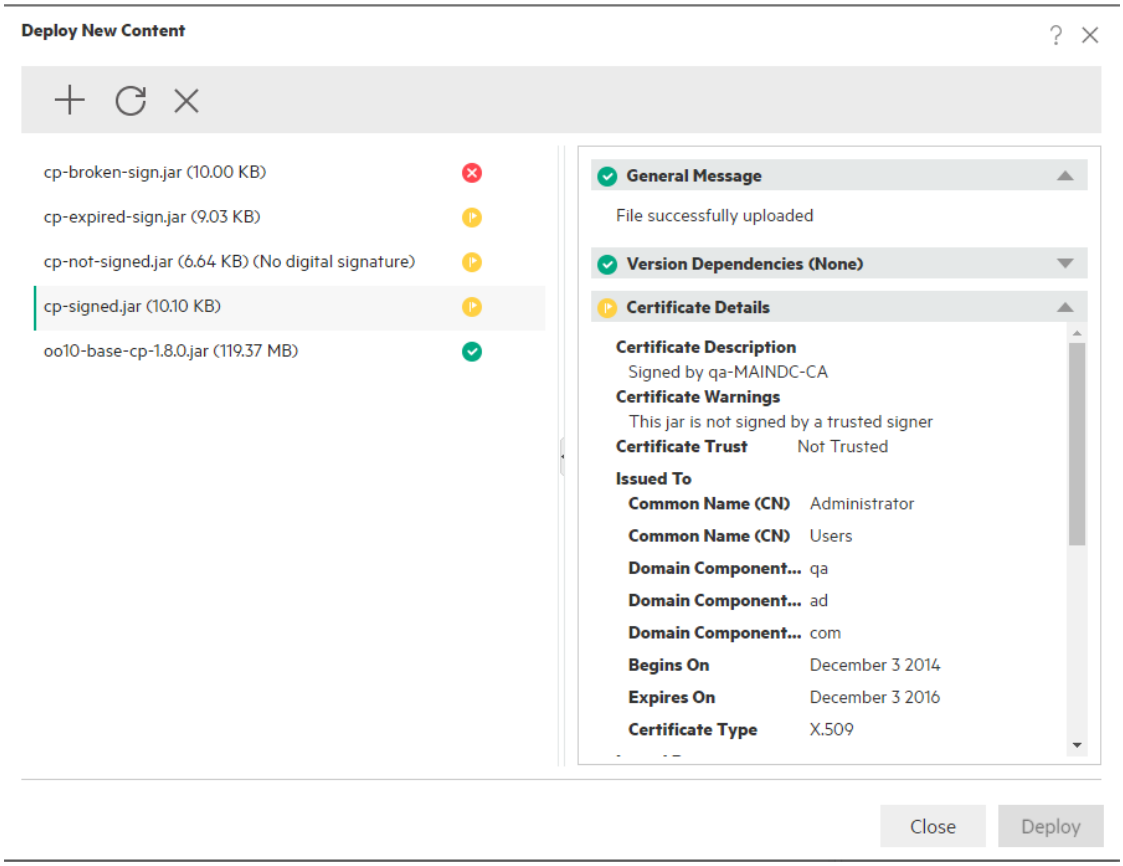

- If a content pack appears with a **Delete** icon, this content pack cannot be deployed. You need to click the **Delete** button in the toolbar to remove this content pack from the dialog box.

  For example, a content pack whose digital signature has been altered cannot be deployed.

- If a content pack appears with a **Warning** icon, this content pack can be deployed, but there is information you need to check about it.

  For example, the signature has expired. In this situation, you need to select the **Ignore expiry of digital signature** check box in order to deploy the content pack.

- If a content pack appears with an **OK** icon, this content pack is okay to be deployed.

- If a content pack has no digital signature, and appear with a **None** icon you can deploy the content pack, but be aware that there is no verification of the content pack.

6. Repeat Steps 3 - 5 to add more content packs to the Deploy New Content dialog box, if desired.
7. Click **Deploy**.

The deployment may take a few minutes. A progress bar displays the progress of the deployment, showing when each step is completed.

> **Note:** It is not possible for two users to deploy a content pack to the same Central at the same time. If you try to deploy content to Central while another user is already deploying content, you will see an error message, indicating that another deployment is currently in progress.

- If deployment was successful, an icon indicating success is displayed in the dialog box and you can click **Close**.
- If deployment fails, messages are displayed, describing the causes of the failure.

  For example, a content pack will fail deployment if it lacks structural integrity (for example, missing references, duplicate UUIDs, duplicate flow paths, duplicate system property names). For more information, see Content Pack Validation - Structural Integrity.

  If the deployment failed, especially in the case of dependency failure, you can click the **Back** button to return to the **Deployment** page. There, you can verify the dependency data, clear the pane, upload all required content packs, and try again.

8. Click **Close** to close the dialog box.

Roll back to the last deployment

After a content pack has been deployed, it is possible to roll back to the last deployment. This might be useful if you find that there are problems with the flows in a new deployment.

If you roll back after a content pack has been deleted, this will undo the deletion.

1. Click the **Content Management**> **Content Packs** tab.
2. Select a content pack and click the **Roll Back Deployment** button.
3. Click **OK** in the confirmation dialog box.

Delete a deployed content pack

1. Click the **Content Management** > **Content Packs** tab.
2. Select a content pack and click the **Delete** button.
3. Click **Delete** in the confirmation dialog box.

Filter content packs

1. Click the **Content Management** > **Content Packs** tab.
2. In the filter box under **CP Name** or **Deployed By**, enter a content pack name, or part of a name, to filter the displayed content packs.

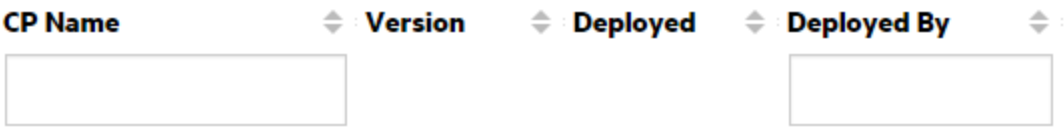

3. To remove the filter, click the **Clear Filters** button.

Display details of a selected content pack

1. Click the **Content Management** > **Content Packs** tab.
2. Select a deployed content pack from the list.
3. Click the **Content** tab in the Details pane, to the right.
4. View the displayed details about the version, publisher, signature, SCM tag, SCM revision (for SVN), description, and the user who deployed it.
5. See how the entities in the content pack are arranged in a tree, under **Library** and **Configuration** folders.
6. Expand and collapse the nodes to see the contents of different folders within the content pack.
7. Clear the check box next to the name of a content pack to remove its details from the Details pane.

Display changes that were made to a selected content pack

1. Click the **Content Management** > **Content Packs** tab.
2. Select a deployed content pack from the list.
3. Click the **Changes** tab in the Details pane, to the right, to see which items have been added, deleted, updated, or moved in the content pack.

   For an item that has been moved, hover the mouse over the item to see where it has been moved from.
4. If desired, you can filter the changes so that only certain types of changes are displayed:
   - Under **Type**, select one or more entity type to display only the changes relating to this type of entity. Options are **Flow**, **Operation**, **Configuration Item**, or **All**.

- Under **Path**, enter part or all of the path to display only changes to entities in a path containing this string.
- Under **Change**, select one or more type of change to display only this type of change. Options are **Added**, **Deleted**, **Modified**, **Moved**, or **Content pack moved**.

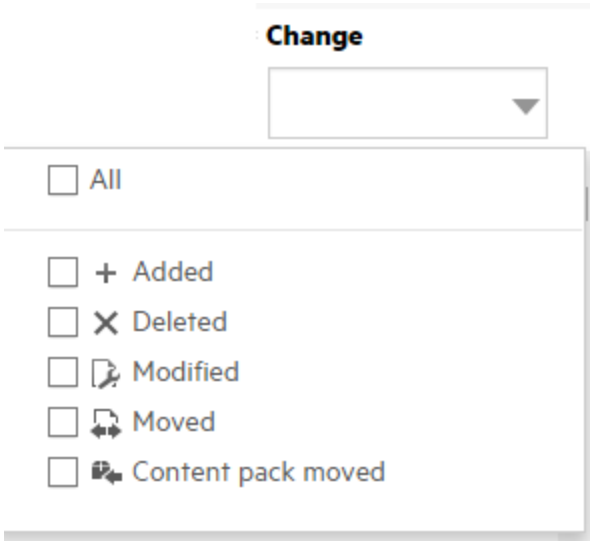

Adjust the display of panes

Use the open/close button in the slider on the edge of the panes to toggle between hiding and displaying the pane.

Drag the slider to adjust the size of a pane.

View statistics about an individual content pack

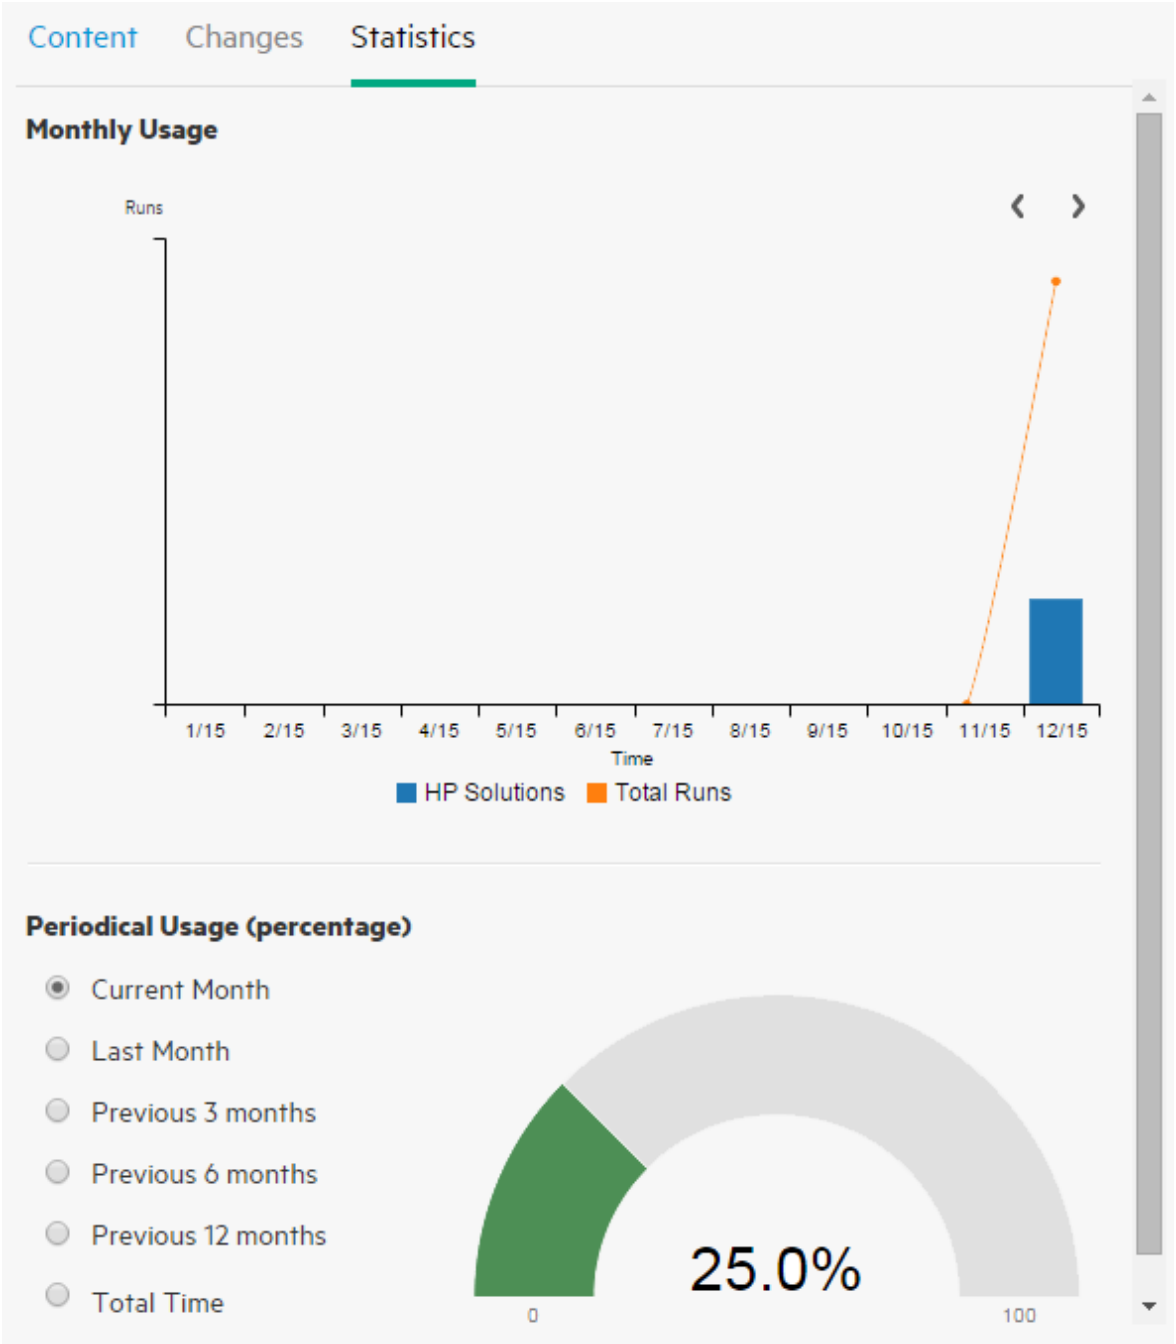

1. Click the **Content Management** > **Content Packs** tab.
2. Select a deployed content pack from the list.
3. Click the **Statistics** tab in the Details pane, to the right.
4. In the **Monthly Usage** section, see how many flows were run were using content from this

deployed content pack for each month since first use, compared with the total runs over the same period.

a. Roll over a month in the bar chart to see a pop-up with a summary of the total number of runs for that month and the number of runs from the content pack.

b. Click the arrows ❮ ❯ to display more months, if not all are displayed onscreen.

c. Click the Legend blocks under the bar chart to filter the display.

5. In the **Periodical Usage (percentage)** section, select a time period, to see how this content pack usage fits into the percentage of all runs during the current month, previous month, quarter, half year, year, or in total.

> **Note:** There may be a delay of 10-20 minutes between the time that a flow is run and when it is reflected in the statistics.

View statistics about all deployed content packs

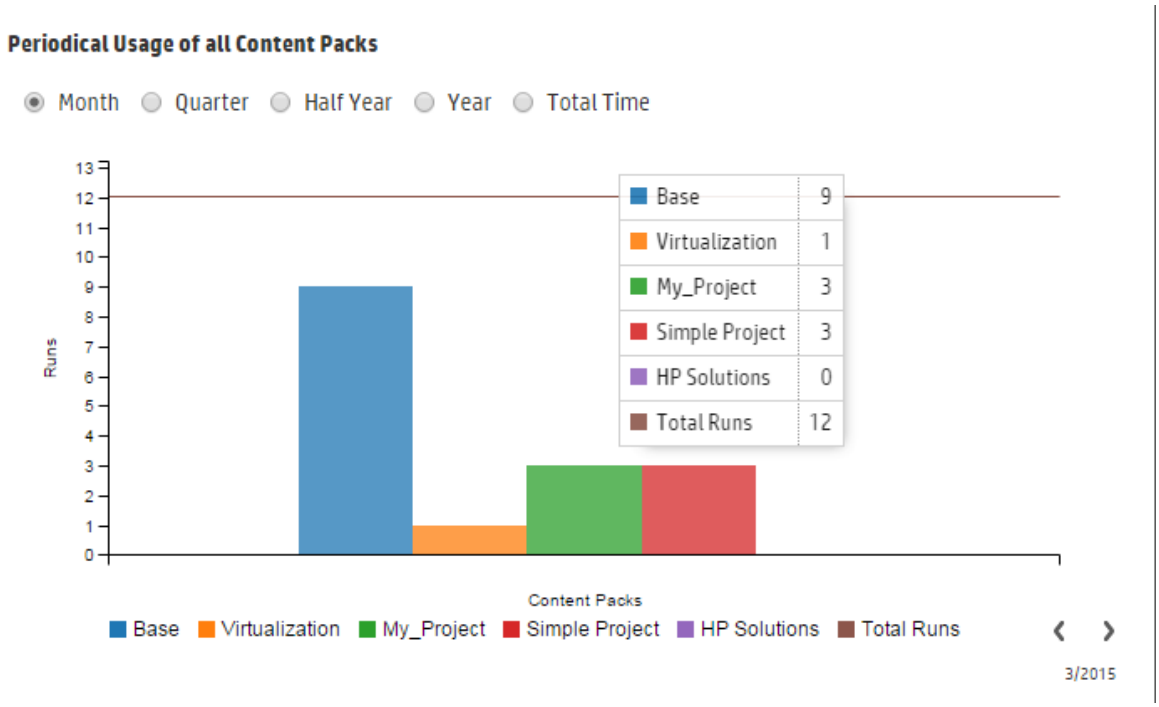

1. Click the **Content Management** > **Content Packs** tab.
2. Click the Statistics button in the toolbar to display the Periodical Usage of all Content Packs dialog box.
3. Select a time period, to compare the number of runs of all deployed content packs over that period.

4. Click the Legend blocks under the bar chart to filter the display.
5. Roll over the bar chart to see a pop-up with a table showing the number of runs for each content pack during the selected period and the total number of runs.
6. Click the arrows ❮ ❯ to display more content packs if not all are displayed onscreen.
7. Click anywhere outside the dialog box to close it.

Reference Material

Deploy New Content dialog box

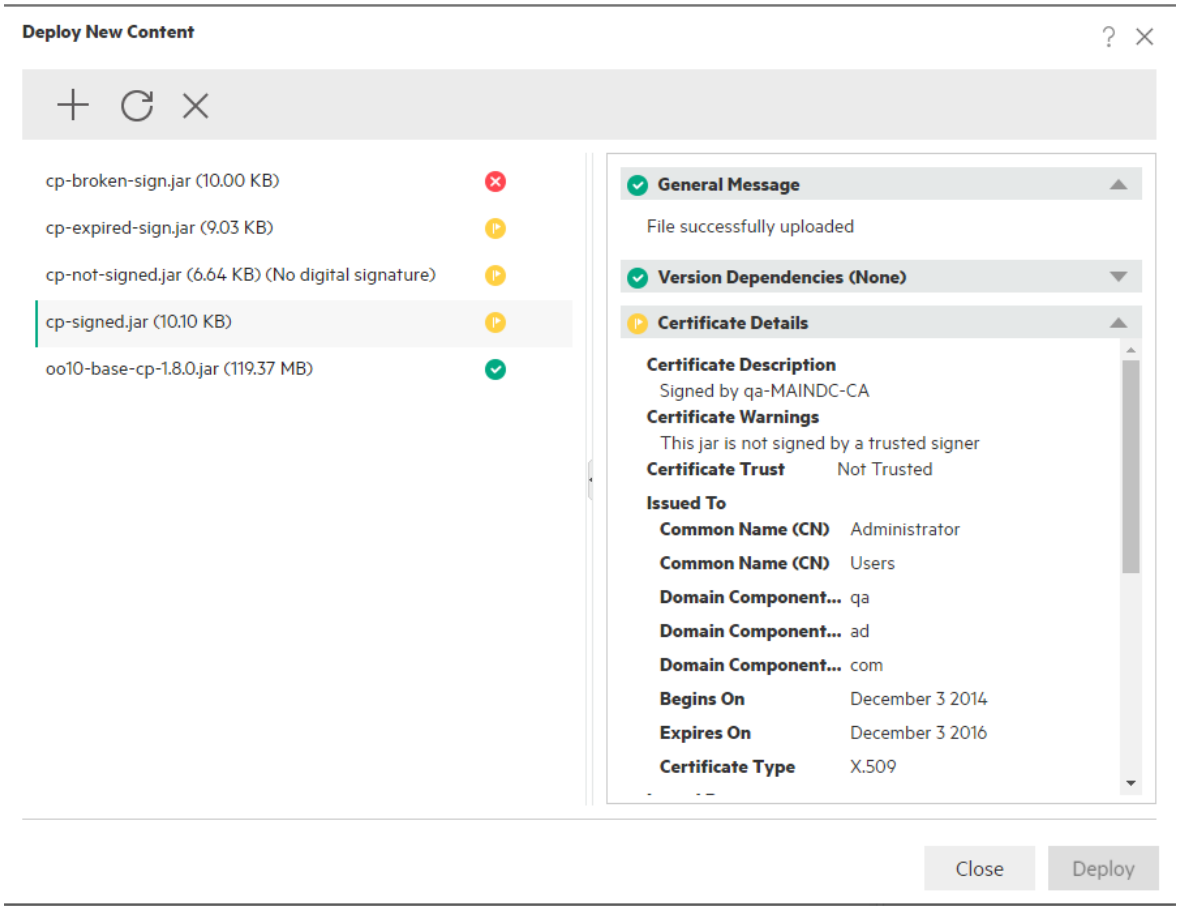

GUI item	Description
General Message	Displays the result of the validation of the content pack during upload. If the content pack was not uploaded to the dialog box, a message explains why.

	For example, if the file is not of the correct file type, or if a content pack with the same name is already uploaded.
Version Dependencies	Displays the result of the dependency validation. If there are missing dependencies, Central displays a warning and displays the details of the content pack versions that contained the missing dependencies, so that you can add them to the deployment. **Name** - the name of the content pack that contained the missing items. **Expected** - the expected version, which originally contained the missing items. This may also be a range of versions (for example, between 1.3 and 1.4, lower than 2.0, higher than 3.0, and so on). **Found** - the existing version of that content pack, which is missing the items.
Certificate Details	Displays the details of the digital certificate: description, trust level, recipient, expiration date, certificate authority, and so on.
Add 	Click to add a new content pack for deployment. You can add multiple content packs and deploy them in a single action.
Remove 	Click to remove the selected content pack from the Deploy New Content dialog box.
Reset 	Click to clear the list of content packs to be deployed, in the Deploy New Content dialog box.
OK 	The content pack is okay to be deployed.
Warning	The content pack can be deployed, but there is information you need to check about it. For example, the signature has expired. In this situation, you need to select the **Ignore expiry of digital signature** check box in order to deploy the content pack.
Delete 	The content pack cannot be deployed. You need to click the **Delete**  button in the toolbar to remove this content pack from the dialog box. For example, a content pack whose digital signature has been altered cannot be deployed.
Deploy	Click to deploy the content packs that are listed in the Deploy New Content dialog box.
Ignore expiry of digital signature	If the digital signature has expired, you need to select this check box in order to deploy the content pack.
Close	Click to close the Deploy New Content dialog box.
Deploy	Click to deploy all verified content packs that are listed in the dialog box.

Content Packs tab

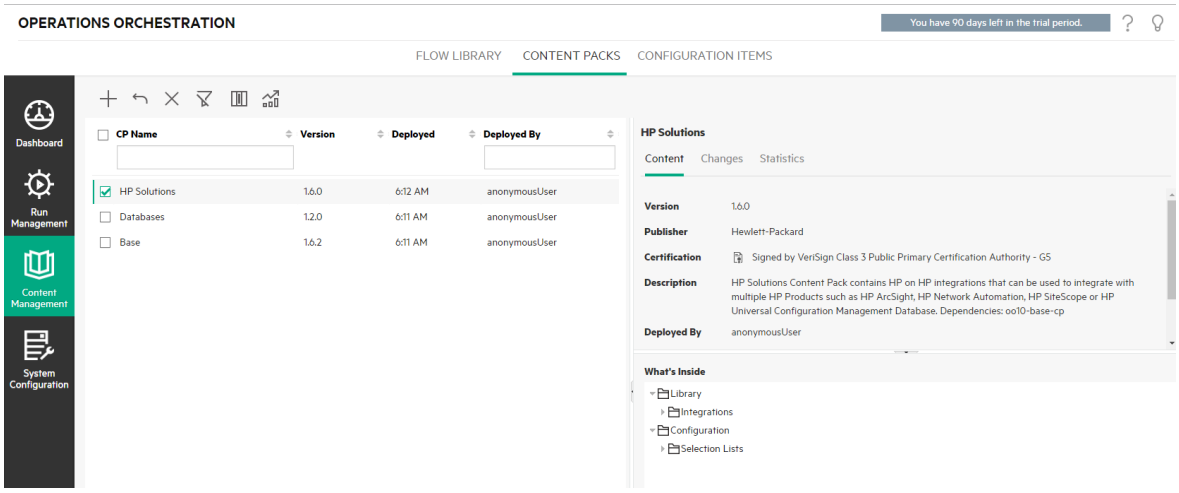

GUI item	Description
Deploy button	Click the **Deploy** button to open the Deploy New Content dialog box.
Roll Back button	Click the **Roll Back** button to roll back to the last deployment.
Delete button	Click to delete the selected content pack.
Clear Filters button	Click to remove the filters and display all the content packs.
Statistics button	Click to display and compare statistical information about all deployed content packs.
Select Columns button	Click to display the column picker, to choose which columns to display.
Clear Filters button	Click to remove the filters and display all the content packs.
Filters	Enter a content pack name or the name of the user who deployed it, or part of a name, to filter the displayed content packs.
Version column	Displays the version number of the content pack.
Deployed column	Displays the time and date when the content pack was deployed.

	Note: If the content pack was deployed during the current day, only the time of day is displayed.
Deployed By column	Displays the user who deployed the content pack.
Content tab	Click the **Content** tab to display the details of the selected content pack. This information includes the: • **Version** • **Publisher** • **Certification**- the digital signature of the content pack, verifying that the content pack was created by a known sender and that it was not altered in transit • **Description** • **Deployed By** - the user who deployed the content pack • **SCM Tag** - from the source control management system in Studio • **SCM Revision** - this is available only for a content pack that was created in a Studio that is connected to SVN Under **What's Inside**, the entities inside the content pack are arranged in a tree, under **Library** and **Configuration** folders, reflecting the way that the content pack appeared in Studio. You can drill down into the folders and see exactly which items exist in the content pack.
Changes tab	Click the **Changes** tab to display the changes that have been made to the selected content pack. This information tells you which items have been added, deleted, updated, or moved. For CloudSlang Content Packs, you can also view Descriptions. For an item that has been moved, hover the mouse over the item to see where it has been moved from.
Statistics tab	Click the **Statistics** tab to display statistical information about the selected content pack.
Open/Close button 	Use the **Open/Close** button in the slider on the left edge of the Details pane to toggle between hiding and displaying the pane.

Content Packs tab > Changes Pane

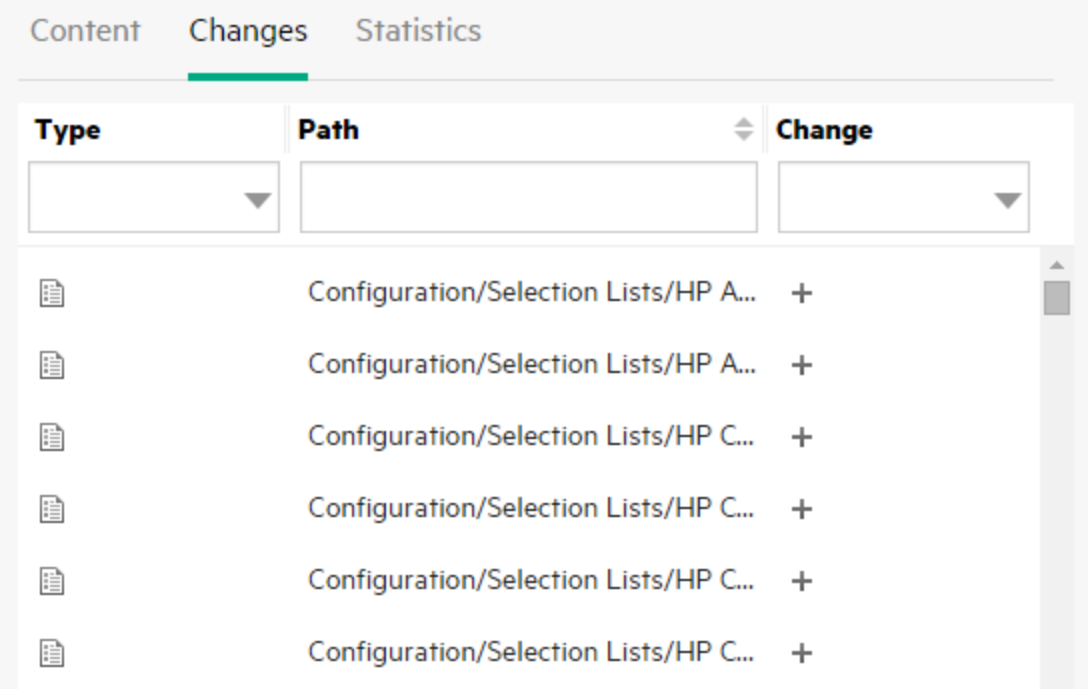

GUI item	Description
Type	Select one or more entity type to display only the changes relating to this type of entity. Options are **Flow**, **Operation**, **Configuration Item**, or **All**.
Path	Enter part or all of the path to display only changes to entities in a path containing this string.
Change	Select one or more type of change to display only changes of this type. Options are **Added**, **Deleted**, **Modified**, **Moved**, or **Content pack moved**.

Content Packs tab > Statistics Pane

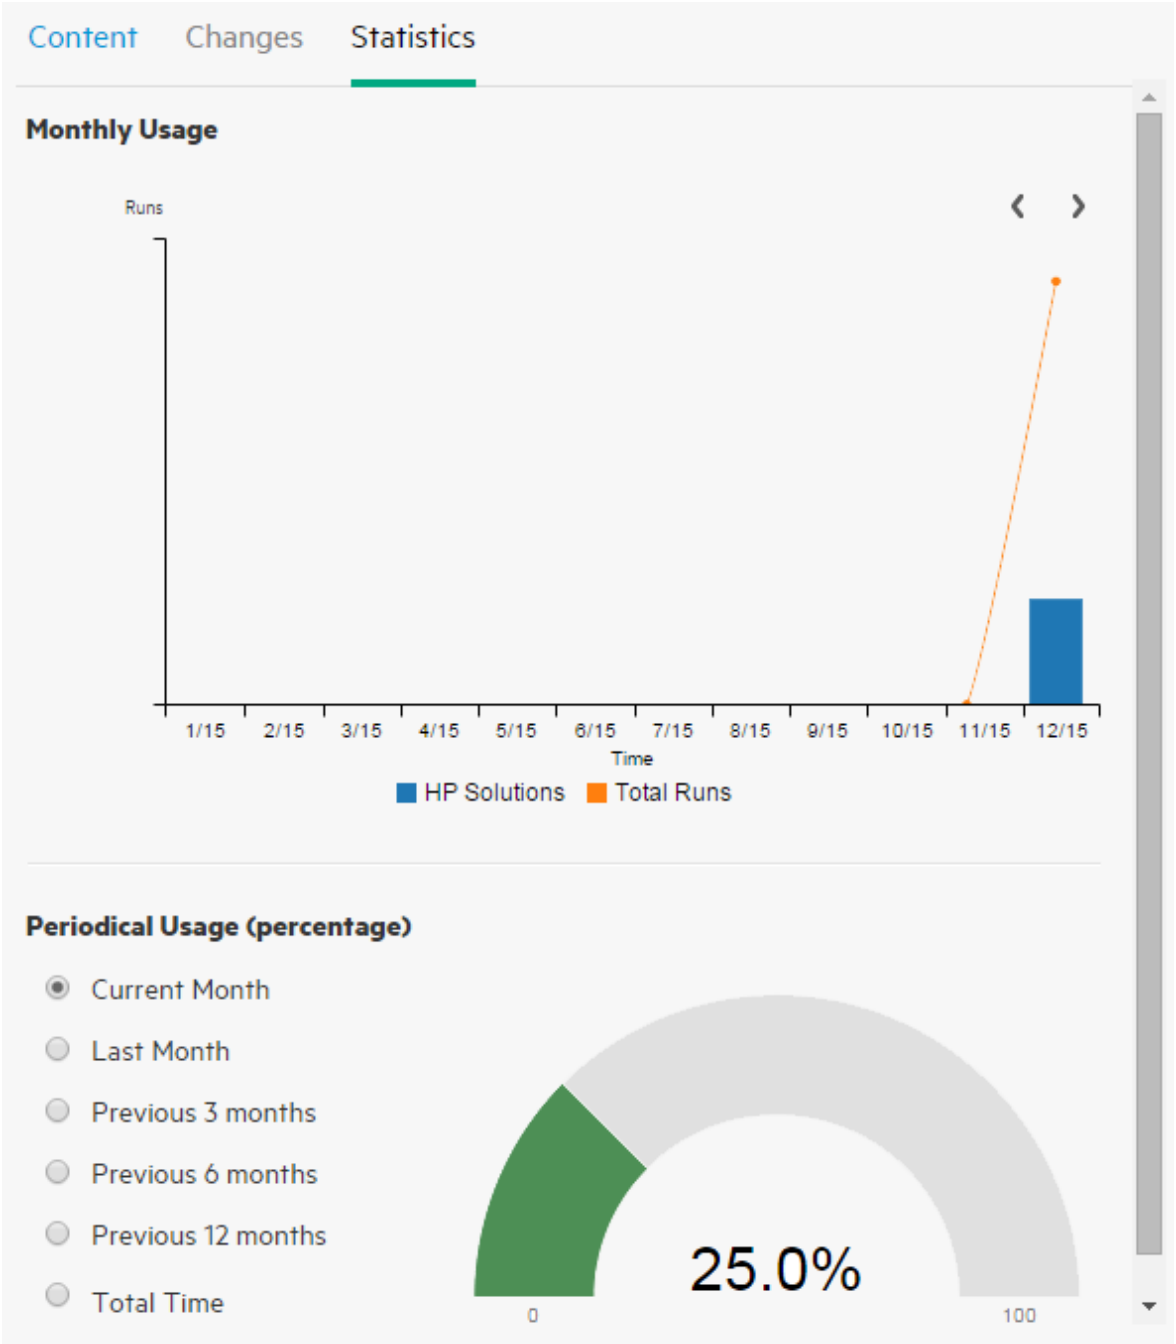

GUI item	Description
Monthly Usage bar chart	See how many flows were run were using content from this deployed content pack for each month since first use, compared with the total runs over the same period. Roll over a month in the bar chart to see a pop-up with a summary of the

	total number of runs for that month and the number of runs from the content pack.
Legend blocks	Click to filter the display.
Arrow buttons ❮ ❯	Click to display more months, if not all are displayed onscreen.
Periodical Usage (percentage)	Select a time period, to see how this content pack usage fits into the percentage of all runs during the current month, previous month, quarter, half year, year, or in total.

Periodical Usage of all Content Packs dialog box

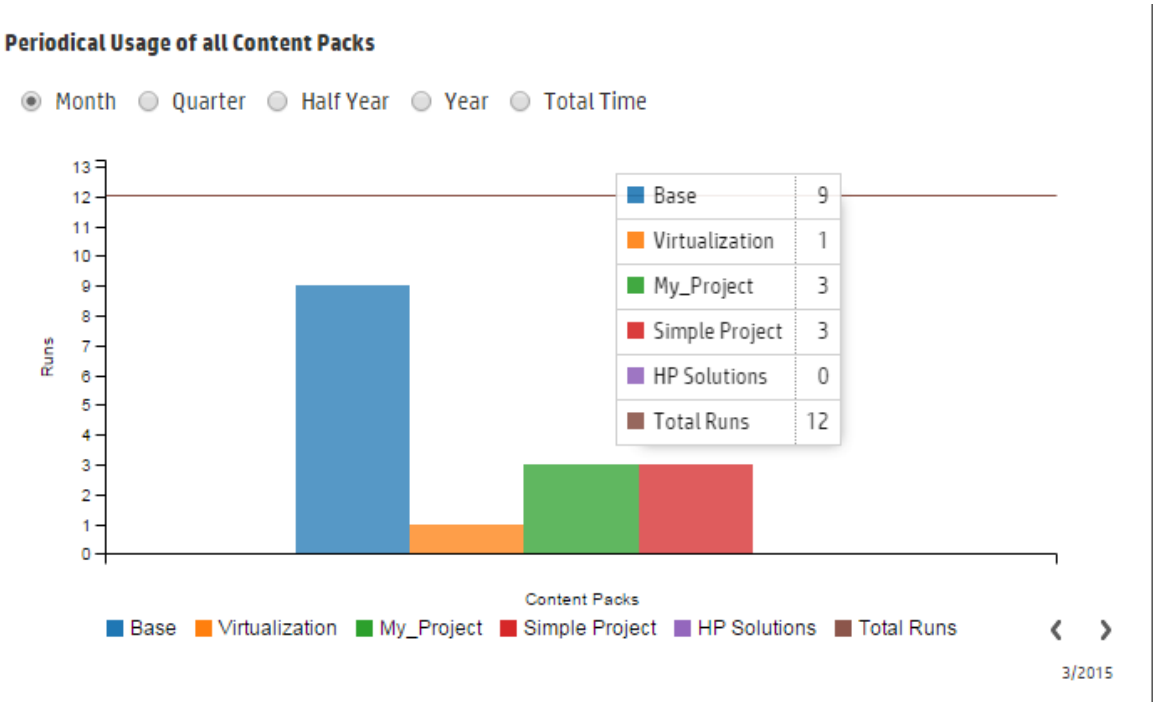

GUI item	Description
Bar chart	Compare the number of flows that were run from the different deployed content packs over the selected period. Roll over the bar chart to see a pop-up with a table showing the number of runs for each content pack during the selected period and the total number of runs.
Periodical Usage options	Select a time period, to compare the number of runs of all deployed content packs over that period.
Legend blocks	Click to filter the display.

Arrow buttons ‹ ›	>Click to display more content packs if not all are displayed onscreen.

Managing the Flow Library

After a content pack has been deployed, you can see the flows inside it in the Flow Library. From here, you can browse or filter flows, to find the one that you need. You can view the flow meta data, view reports about the last time that a flow was run, configure persistence level and run timeout, and set the content permissions.

> **Note:** You will only be able to edit the settings here if you have a role with the **Manage Content Settings** permission.

To access the Flow Library, click **Content Management** > **Flow Library** tab.

> **Note:** The icons of CloudSlang flows and Native OO flows are different.

When you select a flow in the flow library, information about that flow is displayed in the information pane to the right.

This information includes a description of the flow, if one exists, the path to the location where the flow is stored, the flow version, UUID, and ROI.

Configuring the Flow Run Settings for Flows

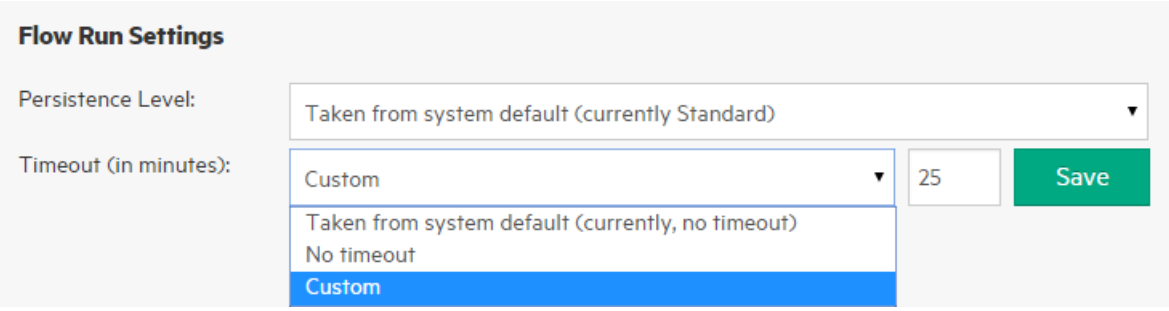

Flow run settings let you control the persistence level and timeout for the selected flow.

- **Persistence Level** defines the level of detail that is saved to the Run Log for the flow. You can choose an extended or standard level, or use the system default. For more information, see Set up the Run Log persistence level for a flow.
- **Timeout** defines how long (in minutes) a flow can run before it is canceled and shows the status

Canceled-Timeout Expired. You can set up a custom timeout setting for the flow, use the system default, or set this flow to have no timeout. For more information, see Set up the timeout for a flow.

- The timeout setting applies to both triggered and scheduled flow runs.
- If you change the timeout value of a flow, this change is applied to new flow runs, but not to those that are already running.
- The timeout setting applied to a flow overrides the system default and is applied even if the system timeout is not enabled.

> **Note:** If a timeout setting was defined for an operation inside a flow, the timeout assigned for the flow is not affected. The flow timeout is triggered according to its value, regardless the values of the timeout values of the steps inside the flow.

Setting Permissions for Content

Content permissions let you restrict access to data (flows and folders) according to the roles that have been assigned to users. This should be done as part of the promotion of a content pack.

You can entitle a role to have either view permission or run permission, or both, for a flow or folder. For example, you might want to give users with the role **Promoter** permission to view and run the contents of a selected folder, and give users with the role **End User** permission to view the content only.

> **Note:** To modify the content permissions, you must be assigned a role with **Manage Content Settings** permission.

If you set the content permissions of a folder and apply the change to all children, these permissions will apply to all the flows within that folder. It is also possible to set the permissions of the entire library.

By default, only roles with permissions are displayed. However, you can select the **Show all roles** check box to display all the roles, including those that do not have permissions for the content.

> **Note:** Content permission is not the same as restricting the actions that HPE OO users are allowed to perform in the system. This is done via Role-Based Access Control (RBAC). For more information about RBAC, see "Setting Up Security – Roles" on page 32

What do you want to do?

Display flow information

1. Click the **Content Management** button to display the Content Management workspace.
2. Click the **Flow Library** tab.
3. Expand the folders to look for the flow that you need.
4. (Optional) Enter part or all of the flow path in the **Filter By** text box to filter the flows.

   For example, if you enter "network", only flows that contain the word "network" or which are located in a path with a folder containing the word "network" are displayed.

   > **Note:** The search is not case-sensitive.

5. Select a flow.
6. View the flow information, displayed in the information pane to the right.
7. If required, use the scroll bar to scroll down and view the contents of the **Description** section.
8. If required, use the sliders at the edges of the information pane to adjust the width and height of the pane.

9. (Optional) Click the **Open Graph** button to display a graphical representation of the flow.

   For more information, see "Viewing a Flow Graph" on page 159.

Display information about the last time a flow was run

You can view the details about what happened the last time a flow was run.

1. Click the **Content Management** button to display the Content Management workspace.
2. Click the **Flow Library** tab.
3. Select a flow in the flow library.

4. Click the **Last Run Report** button. The Drill Down window opens, showing what happened the last time this flow was run.

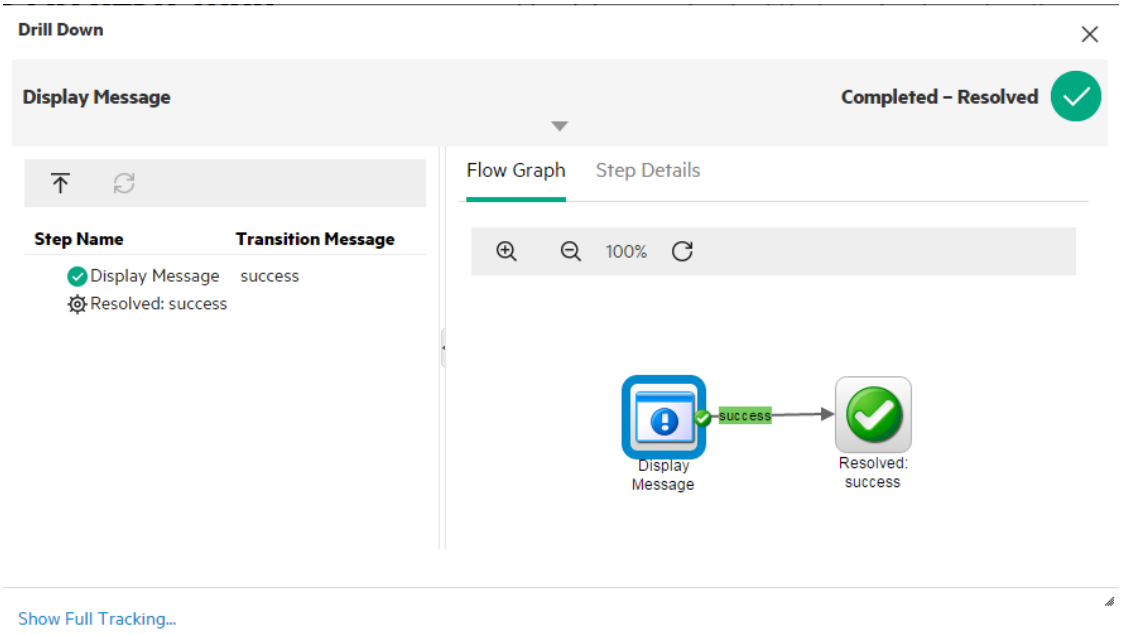

The **Last Run Report** button is only available for flows that have been run at least once in the current session.

The information in the Drill Down window is the same as the information in the Run Explorer drill down view. For more details about this view, see "Testing and Troubleshooting a Flow Run" on page 229.

Set up content permission for a flow or folder

1. Select a flow or folder in the flow library.
2. In the **Permissions** section, click the **Edit** button
3. In the Edit Permissions dialog box, select the **View** and **Run** check boxes to assign permissions for the selected role.

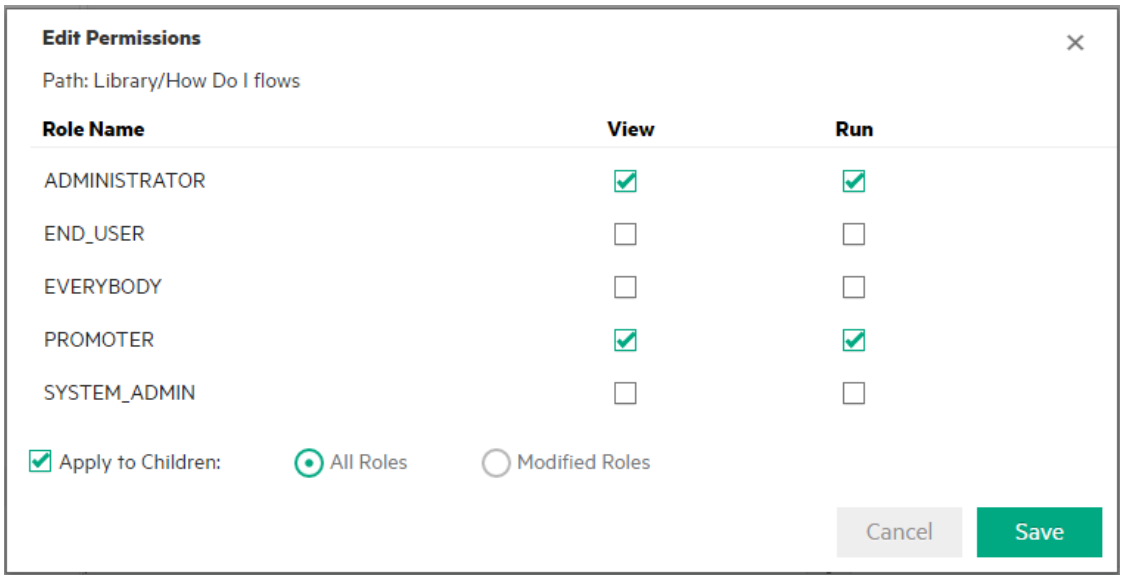

- Select **View** to enable users with the selected role to view this flow or the contents of this folder in Central.
- Select **Run** to enable users with the selected role to run this flow or the contents of this folder in Central.

4. Click **Save**.

Set up the Run Log persistence level for a flow

By selecting the persistence level, you can control the level of detail that is saved to the Run Log for an individual flow.

1. Select a flow in the flow library.
2. From the **Persistence Level** list, select the persistence level for the flow. The options are:
   - **Standard** - Includes raw results, step results, and big inputs/outputs.
   - **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs.
   - **Taken from system default** - whatever was selected as the default persistence level. For information about setting the default persistence level, see "Monitoring and Controlling Database Size" on page 112.

     > **Note:** It is possible to override a flow's persistence level, when running or scheduling a flow.

Set up the timeout for a flow

When you assign a timeout value to a flow, if the flow run exceeds the specified duration, it will be canceled and will show the status **Canceled-Timeout Expired**.

1. Select a flow in the flow library.
2. From the **Timeout (minutes)** list, select the timeout setting for the flow. The options are:
   - **Taken from system default** - whatever was selected as the default timeout setting for the system is applied for this flow. For information about setting the default timeout, see "Setting Up the System Settings - General Settings" on page 100.
   - **No timeout** - no timeout will be applied to this flow, even if a default timeout has been set for the system.
   - **Custom** - lets you specify a timeout for this flow.
3. If you selected **Custom**, enter a number larger than 0, to set the number of minutes for the timeout setting.
4. Click **Save**.

Reference Material

Flow Library pane

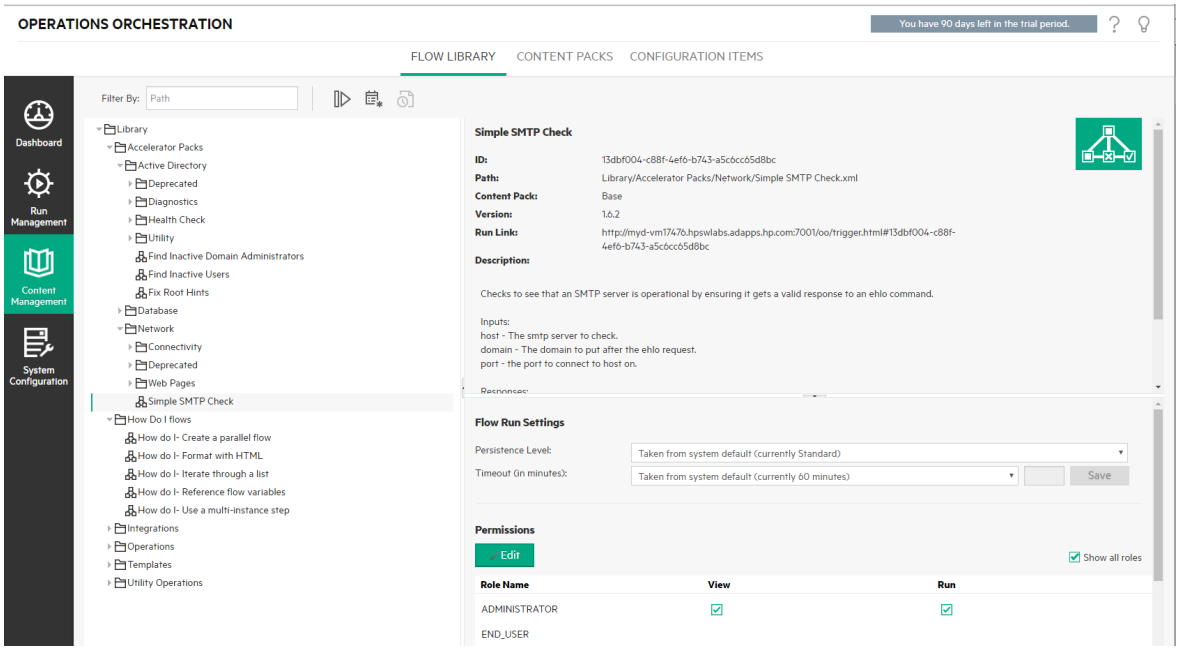

GUI item	Description
Filter By	To locate the flow that you need, enter part or all of the flow path in the **Filter By** text box to filter the flows.
Run  **button**	Click to run the selected flow.
Schedule  **button**	Click to open the Schedule Flow dialog box.
Last Run Report  **button**	Displays the details about what happened the last time that the selected flow was run. The **Last Run Report** button is only available for flows that have been run at least once.
Open Graph Button 	Click the **Open Graph** button to display a graphical representation of the flow. For more information, see "Viewing a Flow Graph" on page 159.
ID	Displays the ID of the selected flow. • Native OO flows IDs have a UUID format. The UUID is unique. • CloudSlang flows, IDs have the `<namepace>.<name>` format. **Note:** The recommended best practice is that the namespace is part of the path and that the name is the filename, but this is not required. The name and namespace are defined inside the yaml file. The id (<namespace>.<name>) must be unique in the system. The filters in the Flow Library use the flow path. So if you don't use the best practice, when you use the filter, using flow name/namespace will not work. For example, in the element `user.flows.hello_world`, the namespace is the folder path **user/flows** and the element name is **hello_world**.
Path	Displays the path to the location of the selected flow.
Content Pack	Displays the content pack that the selected flow came from.
Version	Displays the version number of the selected flow.
Run Link	You can embed parts of the Central UI into an external application. For more information, see "Embedding Central Views in External Web

	Pages" on page 249.
Description	Displays a description of the selected flow. If required, use the scrollbar to scroll down and view the contents of the **Description** section.
Input	Displays the input information for CloudSlang flows only. **Note:** This field may be encrypted as configured by the flow creator.
Output	Displays the output information for CloudSlang flows only. **Note:** This field may be encrypted as configured by the flow creator.
Persistence Level	From the **Persistence Level** list, select the persistence level for the flow. The options are: • **Standard** - Includes raw results, step results, and big inputs/outputs. • **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs. • **Taken from system default** - whatever was selected as the default persistence level. For information about setting the default persistence level, see "Monitoring and Controlling Database Size" on page 112. **Note:** It is possible to override a flow's persistence level, when running or scheduling a flow.
Timeout (minutes)	Enter the timeout setting, so that a flow run will be canceled if it exceeds this duration. The options are: • **Taken from system default** - whatever was selected as the default timeout setting for the system is applied for this flow. For information about setting the default timeout, see "Setting Up the System Settings - General Settings" on page 100. • **No timeout** - no timeout will be applied to this flow, even if a default timeout has been set for the system. • **Custom** - lets you specify a timeout for this flow. Enter a number larger than 0, to set the number of minutes for the timeout setting. **Note:** A custom timeout setting overrides the system default timeout setting and is applied even if the system timeout is not enabled.
Permissions	Displays the permissions that have been set for this flow or folder, for

	each role that has been defined in the system: • Displays **View** if users with the adjacent role have permission to view this flow (or the contents of this folder) in Central. • Displays **Run** if users with the adjacent role have permission to run this flow (or the contents of this folder) in Central. The display of roles in this section corresponds to the roles that have been defined in the system.
Edit	Click to open the Edit Permissions dialog box for the selected flow or folder.
Show all roles	By default, only roles with permissions are displayed. However, you can select the **Show all roles** check box to display all the roles, including those that do not have permissions for the content.

Edit Permissions dialog box

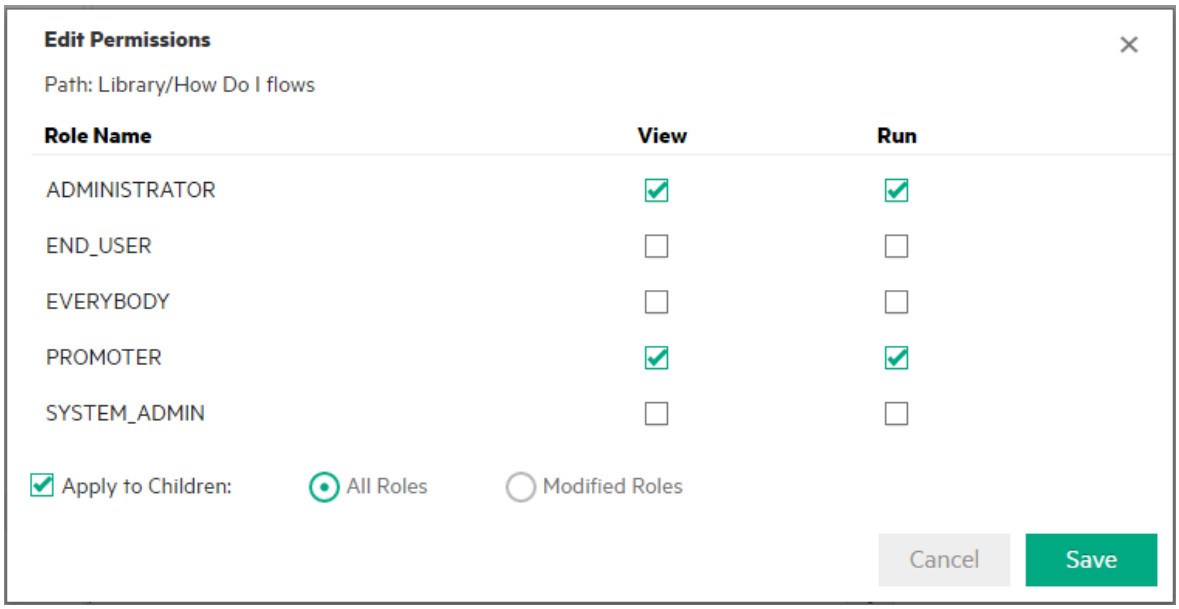

GUI item	Description
View	Select **View** to enable users with the selected role to view this flow (or the contents of this folder) in Central.
Run	Select **Run** to enable users with the selected role to run this flow (or the contents of this folder) in Central.
Apply to children	If you have selected a folder, select the **Apply to Children** check box to apply the updated permissions to all the flows in the folder.

All Roles/Modified Roles	Select whether to apply the changes to all the roles in the folder or to modified roles only.
Save	Click to save the changes to the permissions.

Viewing a Flow Graph

 From the Flow Library or Flow Launcher, you can click the **Open Graph** button to display a flow graph representing the flow. This is the same graph that existed in Studio, when the flow was created.

Note: This feature is not supported when the flow was written in CloudSlang. A "not supported" message is shown when you click the **Open Graph** button.

Why View a Flow Graph?

- In the promotion phase, the Op Admin can view the flow graph in order to better understand what the flow does and which steps and outputs to expect.
- In the execution phase, the End User can use the flow graph to get a clear picture of what the flow does and to verify that this is the flow they need.
- In the troubleshooting phase, the End User and Op Admin can use the flow graph to understand exactly which series of steps caused a flow to fail.

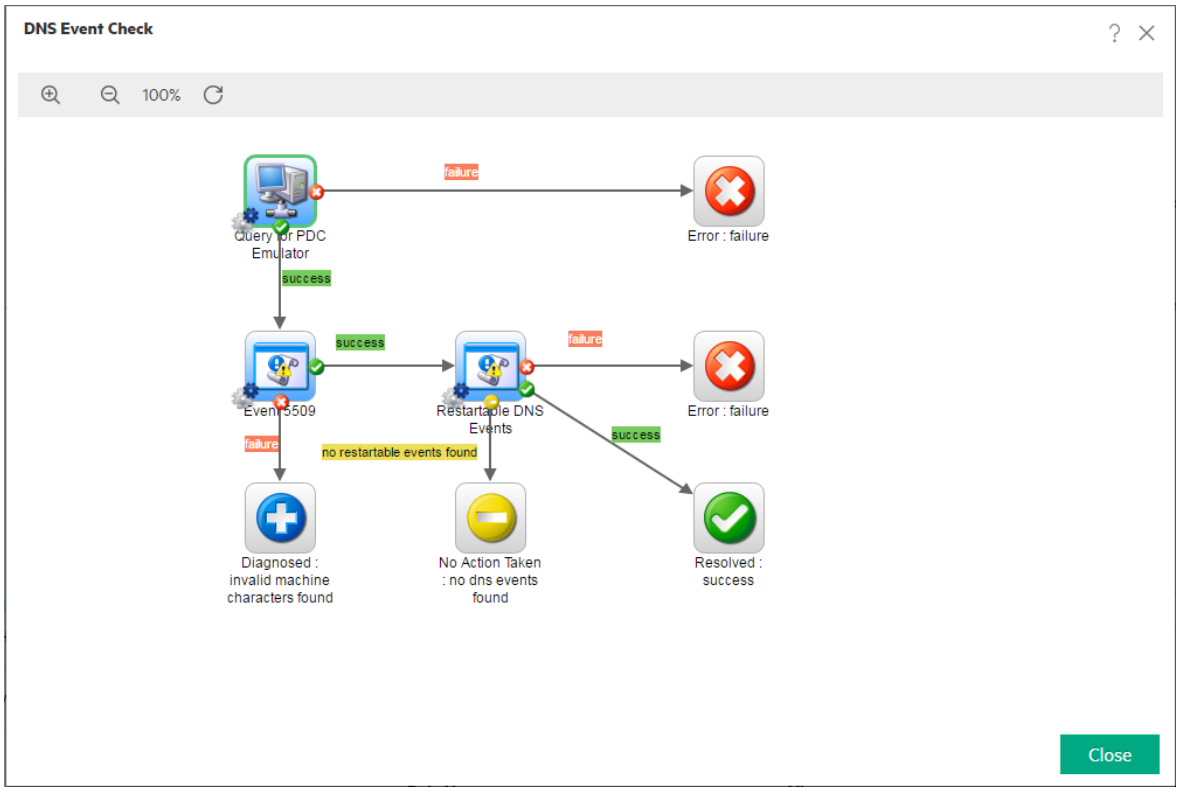

What do you want to do?

Display a flow graph

1. Open either the **Content Management** > **Flow Library** or the **Run Management** > **Flow Launcher**.
2. Select the flow that you want to view.
3. Click the **Open Graph** button  to display a flow graph representing the flow.

Adjust the zoom of a flow graph

When you display a flow graph, the zoom level is automatically adjusted so that the entire flow fits into the window. However, you can increase or decrease this zoom level.

1. Display the flow graph for the selected flow.
2. Adjust the zoom levels:
   - Click the **Zoom In** button to increase the zoom.
   - Click the **Zoom Out** button to decrease the zoom.
   - Click the **Reset** button to reset the view to the default size and position.

Pan across a flow graph

1. Display the flow graph for the selected flow.
2. Roll your mouse over the flow graph so that the cursor changes to a hand icon.
3. Press down on the cursor and drag across the graph to move the display to a specific location on the graph.

Reference Material

Flow graph elements

GUI item	Description
Operation 	An operation has a gray background.
Responses 	Operations have predefined response icons, for “error”,“resolved”, “diagnosed” and “no action”.
Non-blocking 	A non-blocking step does not block the rest of the flow. While it is running, the flow run continues to carry out the steps that come after it. A non-blocking step appears with an orange lightning bolt.

Subflow 	A subflow is a flow within a flow. A subflow appears with a blue background.
Response steps 	Response steps are end points for a flow.
 **Transitions**	Transitions between steps have labels, color-coded using the same colors as the response steps.
Gated transition 	A gated transition restricts access to the next step to users who have been assigned a particular role. Gated transitions appear with a red arrow.
Handoff transition 	A handoff transition is one where a message appears, suggesting that the current user hand off the flow run to another user. Handoff transitions appear with a hollow arrow head.
Multi-instance step 	A multi-instance step can run multiple times, with multiple targets. A multi-instance step resides within a multi-instance branch.
Parallel Split Step 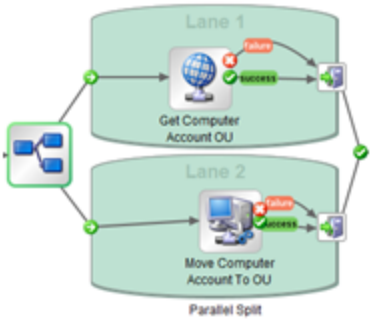	A parallel split step is a set of step sequences that are carried out simultaneously. Each series of steps is represented visually in the flow diagram as a lane.

Flow graph toolbar

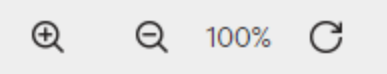

GUI item	Description
Zoom In	Click the **Zoom In** button to increase the zoom.
Zoom Out	Click the **Zoom Out** button to decrease the zoom.
Reset button	Click to reset the view to the default size and position.

Setting Up Configuration Items for a Content Pack

 After a content pack has been deployed to a Central server, the Op Admin may need to configure the contents of this content pack, in order to adjust it to the environment. This includes setting up the:

- **System accounts** - In many cases, content packs come with system accounts configured. For example, if you are deploying a content pack for an Oracle database, it will include Oracle user accounts. You will need to create Central user names and map them to the system accounts in the content pack.
- **System properties** - In many cases, content packs come with system properties configured. You may want to override these properties in Central.

  > **Note:** If a system account or system property was created via API, it is not possible to edit it in the Central UI. You can only edit it via API.
  >
  > In the Central UI, you can edit system accounts and system properties that were created in a content pack in Studio.

- **Group aliases** - If the content pack includes operations that have been assigned to group aliases, rather than actual worker groups, you will need to map these group aliases to actual worker groups. For more information about group aliases, see "Worker Groups and Group Aliases" in the *Concepts Guide*.

  > **Note:** If a group name and alias name are the same, these are mapped automatically by default.

To access the configuration items, click the **Content Management** > **Configuration Items** tab.

> **Note:** The **Configuration Items** tab is only visible if you have been assigned a role with the **Manage Configuration Items** or **View Configuration Items** permission. You will only be able to edit the items if you have a role with the **Manage Configuration Items** permission.

Assigning Permissions to a System Account

You can assign permissions for system accounts. This enables the administrator to control which users can view which system accounts and run (or resume) 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.

- Only users whose roles have **View and Run** permission on a system account will be able to view it in the Configuration Items tree.
- Only users whose roles have **View and Run** permission on a system account will be able to run a flow that uses this system account at the flow level. If a user without permissions tries to run such a flow, it will fail with a “Failed to Complete” status.
- Only users with **View and Run** permission on a system account will be able to remote debug a flow that uses it, from Studio.

You can assign permissions for multiple system accounts at once, by assigning permissions to folders.

If you do not specify the permissions, the system account inherits the permissions from its nearest parent.

When you move a system account, if it has permissions configured, these permission remain configured. If not, the system account inherits the permissions from its nearest parent in the new location.

Viewing Information About Configuration Items

The configuration items are displayed in a tree, with icons representing the different types of item.

Next to the Configuration Items tree, two columns display the values of the items:

- The **Deployed Value** column displays the original value of the configuration item from a deployed content pack.
- The **Central Value** column displays the new value, if the original value was overridden in Central.

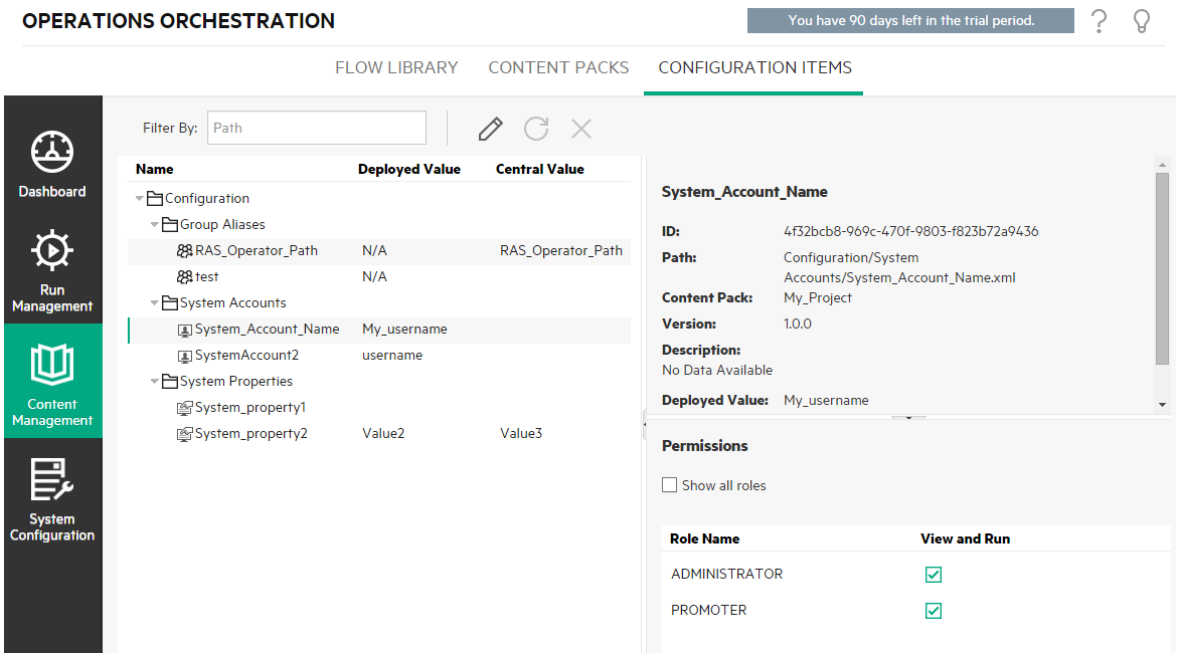

When you select an item in the Configuration Items tree, information about that item is displayed in the information pane to the right. This information includes the ID of the item, the path to the location where the item is stored in the content pack, a description, information about the content pack, and so on.

For system accounts, permissions are also displayed. By default, only roles with permissions are displayed. However, you can select the **Show all roles** check box to display all the roles, including those that do not have permissions for the account.

Deployed Values and Central Values

If you override the value of a configuration item in Central, the **Deployed** value and the **Central** value are connected. If you move or rename a configuration item in Studio and deploy it again, the **Central** value is also moved or renamed.

Path Uniqueness for Configuration Items

Configuration items must be unique. If you deploy a configuration item in a folder containing a configuration item with the same name, it will either overwrite that item (if they have the same UUID) or fail deployment (if it has a different UUID).

> **Note:** However, it is possible to have two configuration items with the same name if they are located in different folders.

There are limitations to the way that Central enforces path uniqueness:

- Unlike Studio, Central enforces path uniqueness for certain item types only – namely, flows and the specific configuration item types that are supported by the REST API **/v1/config-items**. For more information, see the *HPE OO API Guide*.
- Unlike in Studio, paths are case-sensitive in Central. This means that Central may not detect a conflict between “my item” and “My Item”, for example.

System Accounts in Resumed Flows

When you resume a flow, HPE OO fetches the system account configuration from the current system. So if a system account was changed between the time the flow was started and the time it was resumed, HPE OO takes the new system account value.

What do you want to do?

Display configuration items in a content pack

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. In the Configuration Items tree, expand the folders to look for the item that you need.
4. (Optional) Enter part or all of the configuration item path in the **Filter By** text box to filter the items.

   For example, if you enter "network", only items that contain the word "network" or which are located in a path with a folder containing the word "network" are displayed.

   > **Note:** The search is not case-sensitive.

5. Select a configuration item or folder and view the information displayed in the pane to the right.
6. If required, use the scroll bar to scroll down and view the contents of the **Description** section.
7. If required, use the sliders at the edges of the information pane to adjust the width and height of the pane.

8. In the case of a system account, you can also view the roles that are assigned **View and Run** permission for the selected folder or account. By default, only roles with permissions are displayed. However, you can select the **Show all roles** check box to display all the roles, including those that do not have permissions for the account.

> **Note:** To add **View and Run** permission for a system account or folder, click the **Edit** button.

Add configuration items to a content pack

It is not possible to add configuration items to a content pack in Central. These should be created in Studio and deployed to Central.

Edit a system account in a content pack

If a system account was created via API, it is not possible to edit it in the Central UI. You can only edit it via API.

In the Central UI, you can edit system accounts that were created in a Studio content pack.

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. Select the system account that you want to edit and click the **Edit** button.

   The Edit System Account dialog box opens. This dialog box will vary depending on the permissions of the currently logged in user. For example, if the logged in user is not entitled to assign permissions, the **Permissions** section is not shown in the dialog box.

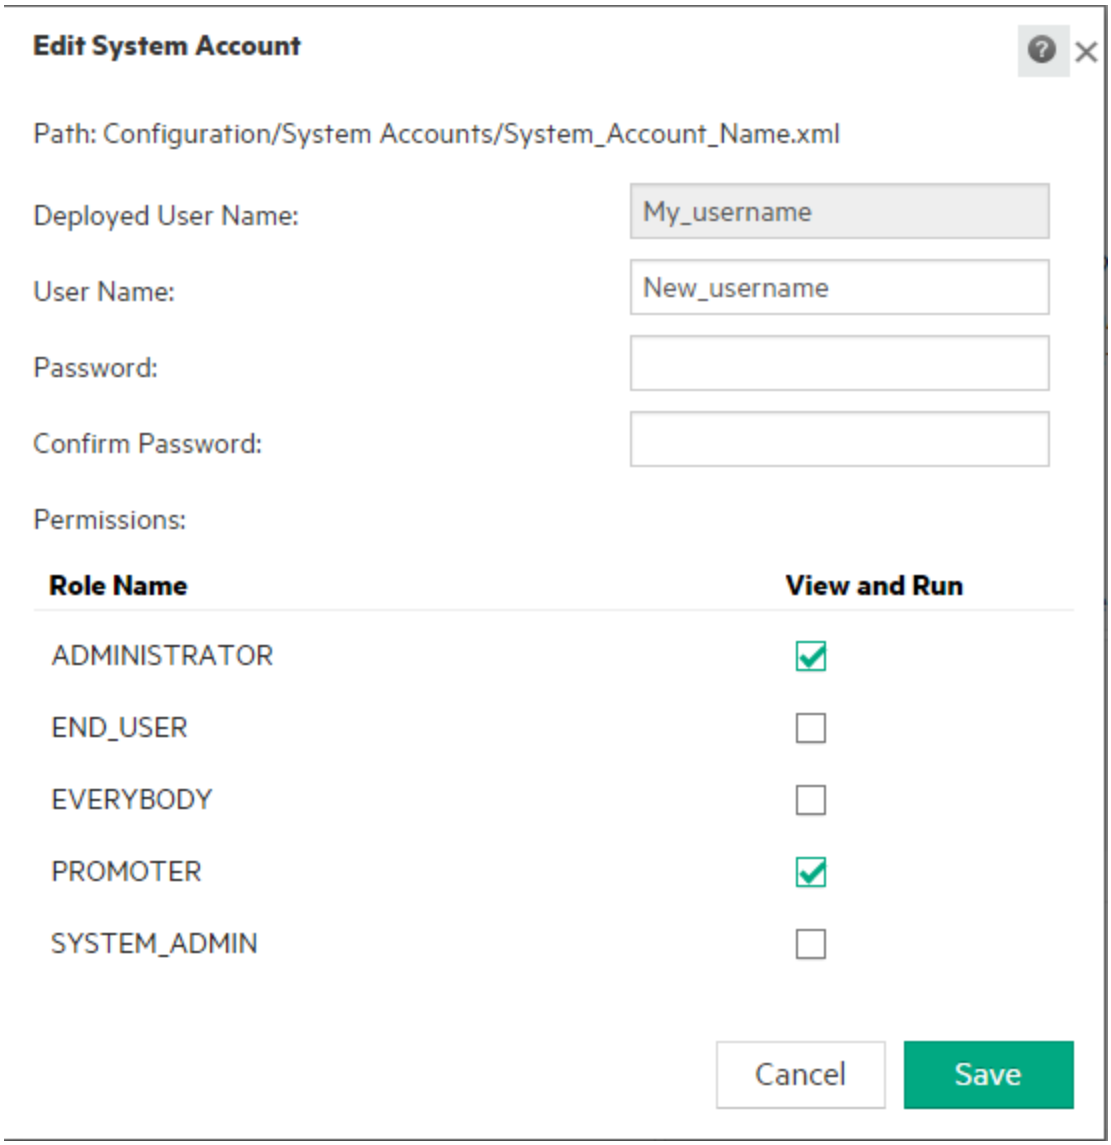

The **Deployed User Name** box displays the system account user name as it appears in the content pack, as defined in Studio.

4. In the **User Name** box, enter the Central value for the system account user name. This user name will be mapped to the system account from the content pack, overriding the deployed user name.

> **Note:** The user name does not have to be identical to the deployed user name.

5. In the **Password** box, enter a password for the user.

> **Note:** If the **Include system account passwords** check box was selected when this content pack was created in Studio, the password is automatically loaded into the **Password**

box.

6. Enter the password a second time in the **Confirm Password** box.
7. In the **Permissions** section, select the check box under **View and Run** for all the roles for which you want to assign permission for this system account.

   Only users with **View and Run** permission on a system account will be able to view it in the Configuration Items tree or to run or resume flows and operations that use this system account.
8. Click **Save**.

Edit permissions for a system account folder

You can assign permissions to a system account folder, which assigns those permissions to all the system accounts in that folder.

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. Select a folder in the Configuration Items tree and click the **Edit**  button.
4. In the **Permissions** section, select the check box under **View and Run** for all the roles for which you want to assign permission for the system accounts in this folder.

   Only users with **View and Run** permission on these system accounts will be able to view them in the Configuration Items tree or to run flows and operations that use these system accounts.

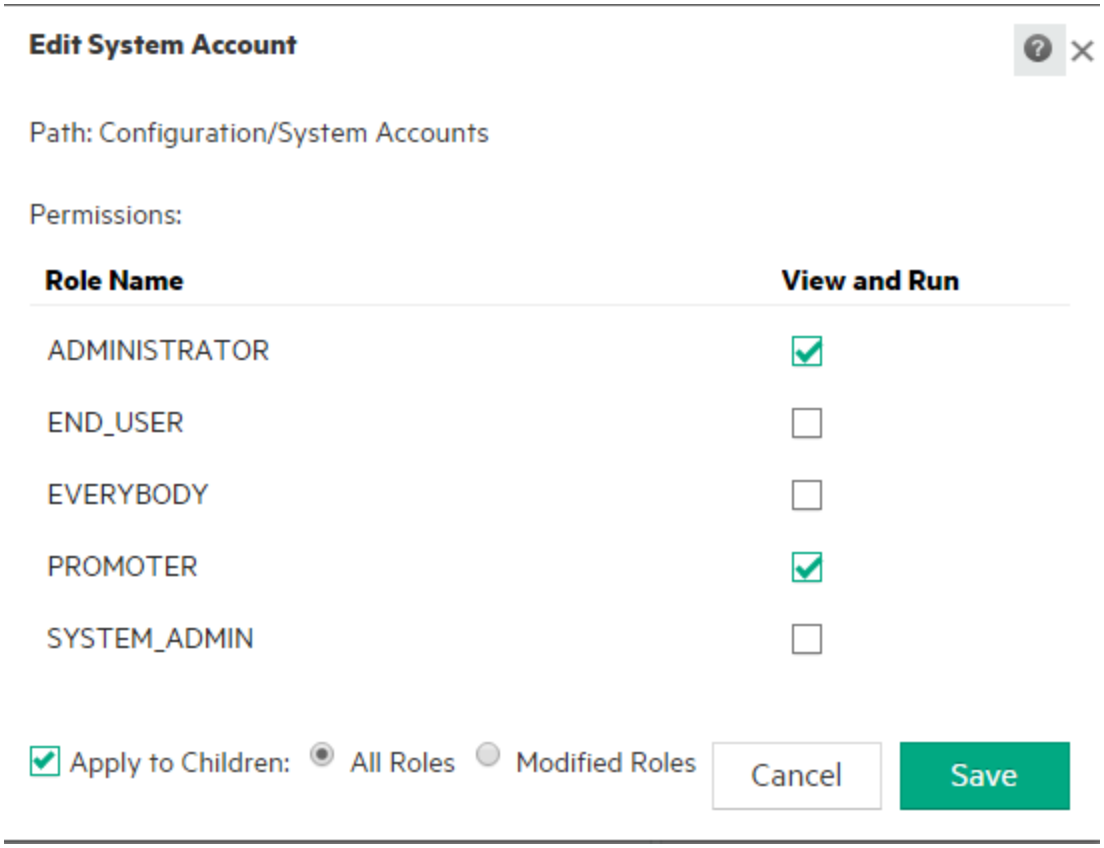

5. Select the **Apply to Children** check box to apply these changes to all the system accounts in the folder.
   - Select **All Roles** to apply the changes to all roles, including those that have not been modified.
   - Select **Modified Roles** to apply the changes only to roles that have been modified.
6. Click **Save**.

Edit a system property in a content pack

If a system property was created via API, it is not possible to edit it in the Central UI. You can only edit it via API.

In the Central UI, you can edit system properties that were created in a Studio content pack.

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. Select the system property that you want to edit and click the **Edit** button.

   The System Property Details dialog box opens.

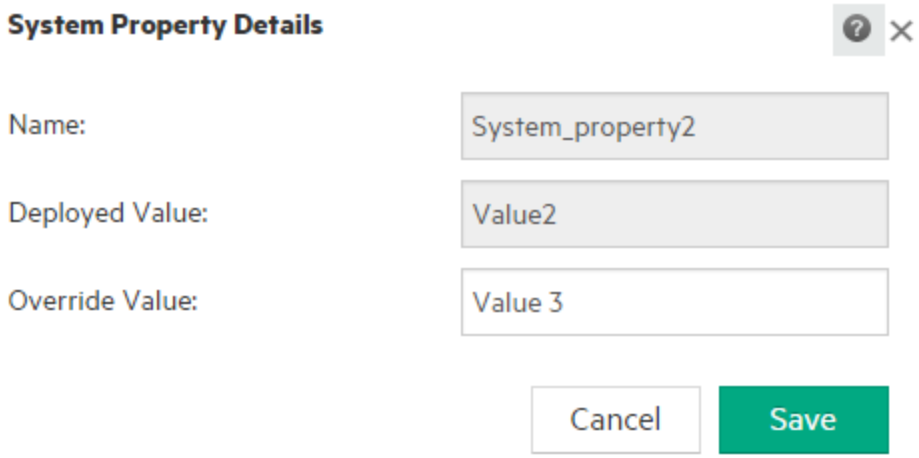

4. In the **Override Value** box, enter a new value for the system property.
5. Click **Save**.

Assign a worker group to a worker group alias

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. Select the worker group alias that you want to edit and click the **Edit** button.

   The Group Alias Details dialog box opens.

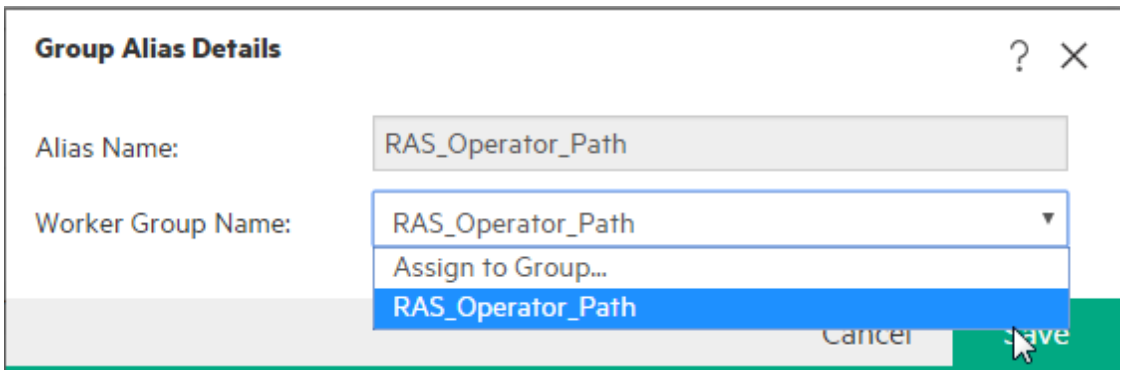

   The **Alias Name** box displays the name of the alias that was used in the operations.
4. From the **Worker Group Name** list, select the worker group to which you want to map this group alias.

5. Click **Save**.

> **Note:** If you want to map the group alias to a group that does not exist yet, go to **System Configuration** > **Topology** > **Workers** and create a new group, and then return to this task to do the mapping. For more information about creating worker groups, see "Setting Up Topology – Workers and RASes" on page 81.

Revert configuration items to their original value

If someone has changed the value of a configuration item that was in the content pack at the time of deployment, you can revert it to the original value.

1. Select the item and click the **Revert**  button.
2. Click **Yes** in the confirmation dialog box.

Delete configuration items in a content pack

It is possible to delete configuration items that have a **Central** value but not a **Deployed** value. For example, if configuration items were deployed prior to HPE OO 10.20, and then HPE OO was upgraded, these will have a **Central** value and no **Deployed** value.

It is not possible to delete a configuration item that is part of a content pack that is currently deployed to Central.

1. Click the **Content Management** button to display the Content Management workspace.
2. Select **Configuration Items**.
3. Select the configuration item that you want to delete and click the **Delete**  button.
4. Click **Yes** in the confirmation dialog box.

Reference Material

Configuration Items Tree

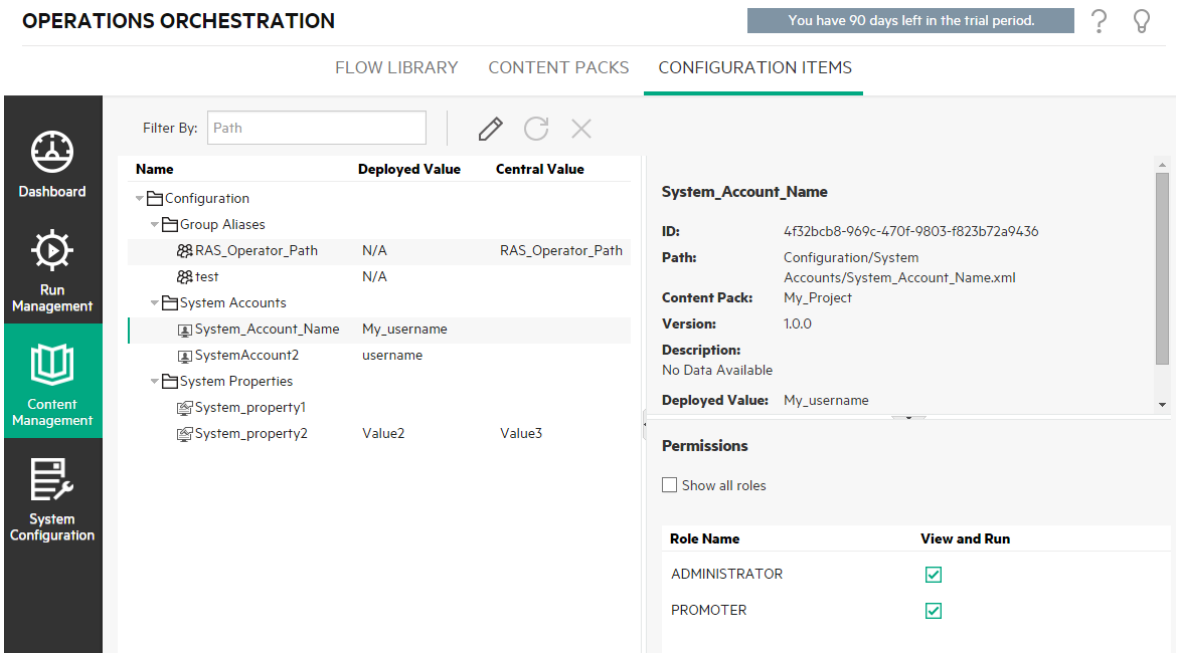

GUI item	Description
Name	Displays the name of the configuration item as it appears in the content pack.
Deployed Value	Displays the original value as deployed to Central.
Central Value	Displays the new value, which overrides the original value.
Edit  **button**	Depending on which configuration item is selected, opens the relevant dialog box, so that you can edit the configuration item.
Revert  **button**	Reverts the selected configuration item to its original deployed value, which was overridden in Central.
Delete  **button**	Deletes the selected configuration item. It is possible to delete configuration items that have a **Central** value but not a **Deployed** value. For example, configuration items that were deployed prior to version 10.02, and HPE OO was upgraded. It is not possible to delete a configuration item that is part of a content pack that is deployed to Central.

Edit System Account dialog box

The Edit System Account dialog box will vary depending on the permissions of the currently logged in user. For example, if the logged in user is not entitled to assign permissions, the **Permissions** section is not shown in the dialog box.

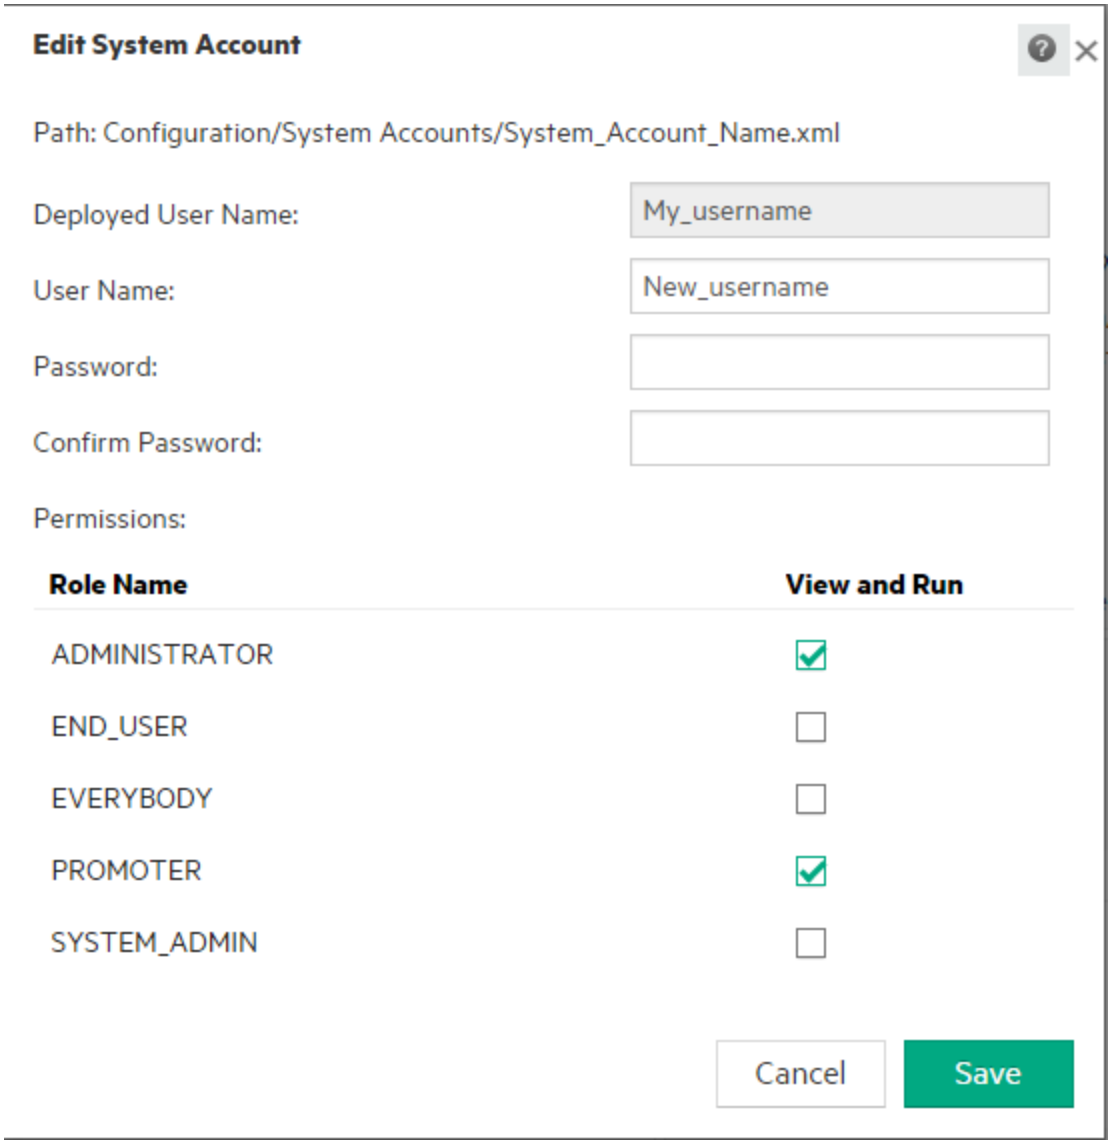

GUI item	Description
Path	Displays the path where the system account is located.
Deployed User Name	Displays the user name of the system account, as defined in the content pack.
User Name	Enter the Central value for the system account user name. This user name

	will be mapped to the system account from the content pack, overriding the deployed user name.
Password	Enter the password for the user name. **Note:** If the **Include system account passwords** check box was selected when this content pack was created in Studio, the password is automatically loaded into the **Password** box. You can change this value.
Confirm Password	Enter the password a second time, for confirmation.
View and Run	Select the check box under **View and Run** for all the roles for which you want to assign permission for this system account. Only users with **View and Run** permission on a system account will be able to view it in the Configuration Items tree or to run flows and operations that use this system account.

Edit System Account dialog box - Folder

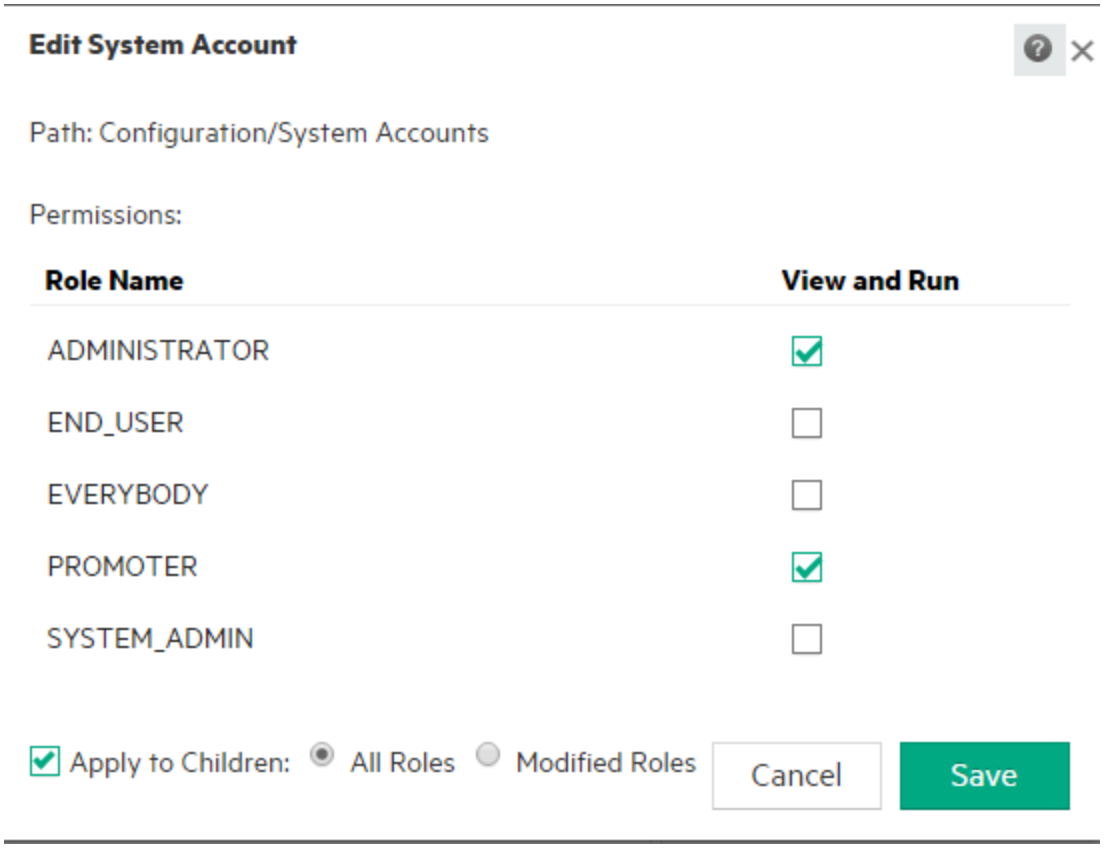

GUI item	Description

View and Run	Select the check box under **View and Run** for all the roles for which you want to assign permission for the system accounts in this folder. Only users with **View and Run** permission on these system accounts will be able to view them in the Configuration Items tree or to run flows and operations that use these system accounts.
Apply to Children: All Roles	Select the **Apply to Children** check box and the **All Roles** option to apply these changes to all the system accounts in the folder, for all roles, including those that have not been modified.
Apply to Children: Modified Roles	Select the **Apply to Children** check box and the **Modified Roles** option to apply these changes to all the system accounts in the folder, but only for roles that have been modified.

System Property Details dialog box

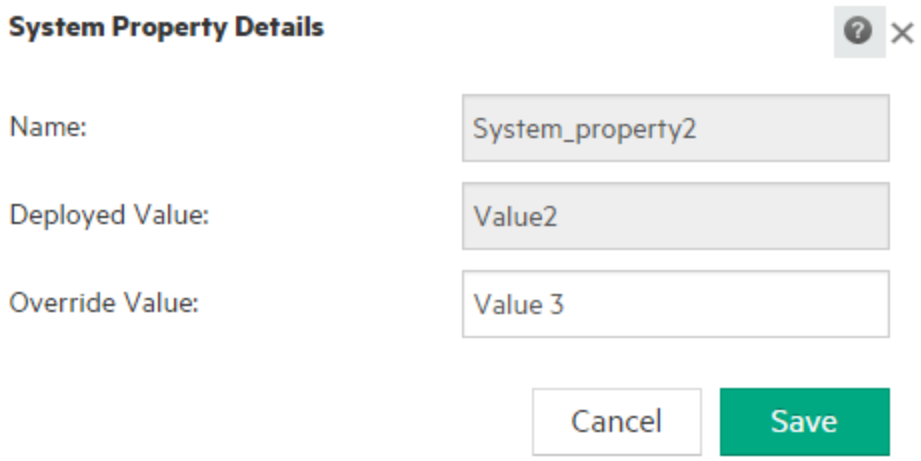

GUI item	Description
Name	Displays the name of the system property, as defined in the content pack.
Deployed Value	If the system property had an original value in the deployed content pack, displays that value.
Override Value	Enter a value for the system property. If the system property had an original value in the deployed content pack, this value will override it.

Group Alias Details dialog box

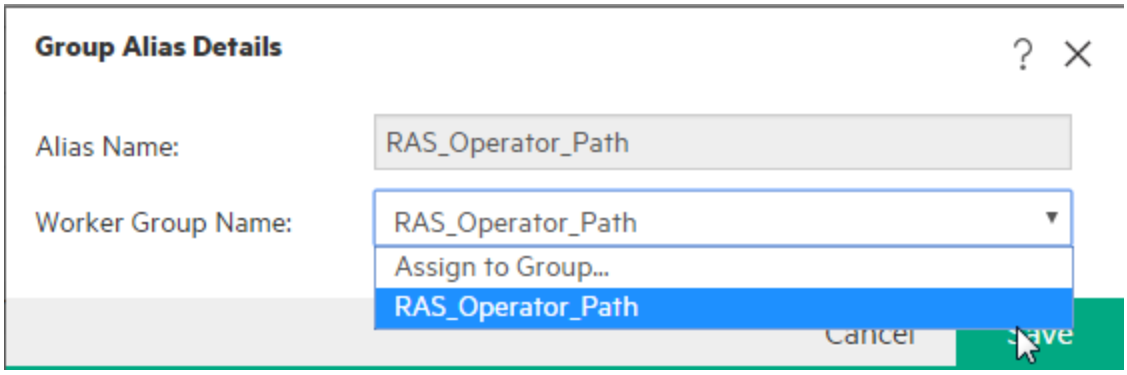

GUI item	Description
Alias Name	Displays the name of the group alias as it appears in the content pack.
Worker Group Name	Select the worker group to which you want to map this group alias.

Testing and Troubleshooting a Content Pack

In this step, the Op Admin tests and troubleshoots the flows in the Staging server.

What do you want to do?

Test a content pack on the Staging server

1. Click the **Content Management** button to display the Content Management workspace.
2. Click the **Flow Library** tab.
3. Select the flow that you want to test, and click the **Run** button.
4. Monitor the progress of the run, and check whether it runs as expected.
5. To further test the flow, run it again, with the same or different parameters.

Running and Monitoring Flows

The Op Admin and the End User are responsible for running and monitoring flows.

Running and Monitoring a Flow – Overview

Once the content packs have been deployed, the Op Admin or End User can run flows and monitor the flow runs.

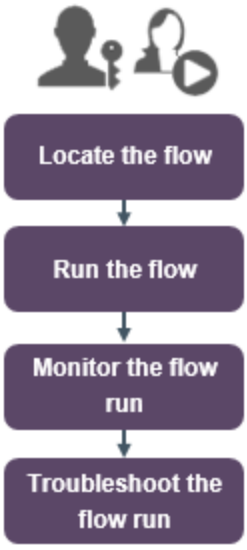

Note that this is just a high level look at the workflow, and there are many options that are not described here. For more detailed information about any of the steps, use the links to learn about the options in-depth.

Step 1: Find the flow that you want to run

 The Op Admin/End User locates the flow either from the Flow Library or the Flow Launcher. Go to the **Content Management > Flow Library** module or to the **Run Management > Flow Launcher** module. See "Running a Flow" on page 181.

Step 2: Run the flow

 The Op Admin/End User runs the flow. See "Running a Flow" on page 181.

 Alternatively, the Op Admin/End User schedules the flow to run later. See "Scheduling Flow Runs" on page 196.

Step 3: Monitor the flow run

 The Op Admin/End User tracks the flow run in the Run Explorer. See "Tracking and Managing Flow Runs" on page 215.

 If required, the Op Admin/End User performs actions on flow runs, such as pausing, resuming, and canceling flow runs.. See "Tracking and Managing Flow Runs" on page 215.

Step 4: Troubleshoot the flow run

 If required, the Op Admin investigates any problems with the run. See "Testing and Troubleshooting a Flow Run" on page 229.

Running a Flow

In the **Flow Launcher** tab in the Run Management workspace, you can select a flow, view its details, prepare it for running, and start the run.

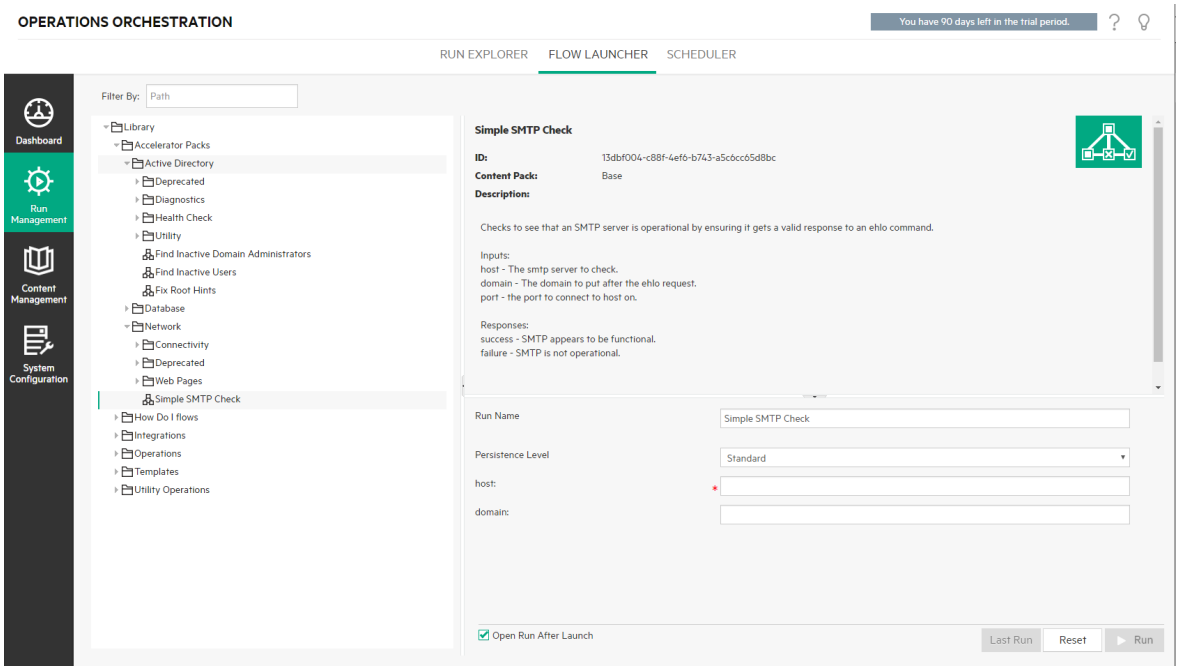

- The flows are displayed in a tree, where you can browse to locate the one you need. You can enter text in the **Filter By** box to filter the display of flows, and you can expand and collapse the tree nodes.
- In the **Flow Details** pane, on the upper right, you can view information about the selected flow, such as ID, content pack, description, and inputs.
- In the **Inputs** pane, on the lower right, you can give the run a name and enter any required inputs.

> **Note:** The flow level input appears in Central only if it can be modified by a user. For example:
>
> - The input was defined as one that prompts the user
> - The value of **Assign From** is the same as the flow input name
>
> For example, the input will not appear if the flow author has not assigned a value to **Assign**

> **From** and has set **Otherwise** to **Use Constant**.
>
> For more information see “Input Inspector > Inputs tab” in the *HPE OO Studio Authoring Guide*.

When you have entered the inputs and the flow is ready, you can click the **Run** button to start the flow.

> **Note:** It is also possible to run a flow from the Flow Library (in **Content Management**).

If you have run this flow previously, you can click the **Last Run** button to view the progress of your last run in the Drill Down dialog box.

For information about how to pause, resume, or cancel a flow run, see "Tracking and Managing Flow Runs" on page 215.

What do you want to do?

Run a flow

1. Click the **Run Management** button.
2. Click the **Flow Launcher** tab.
3. In the Flows tree on the left, browse to find the flow that you want to run.

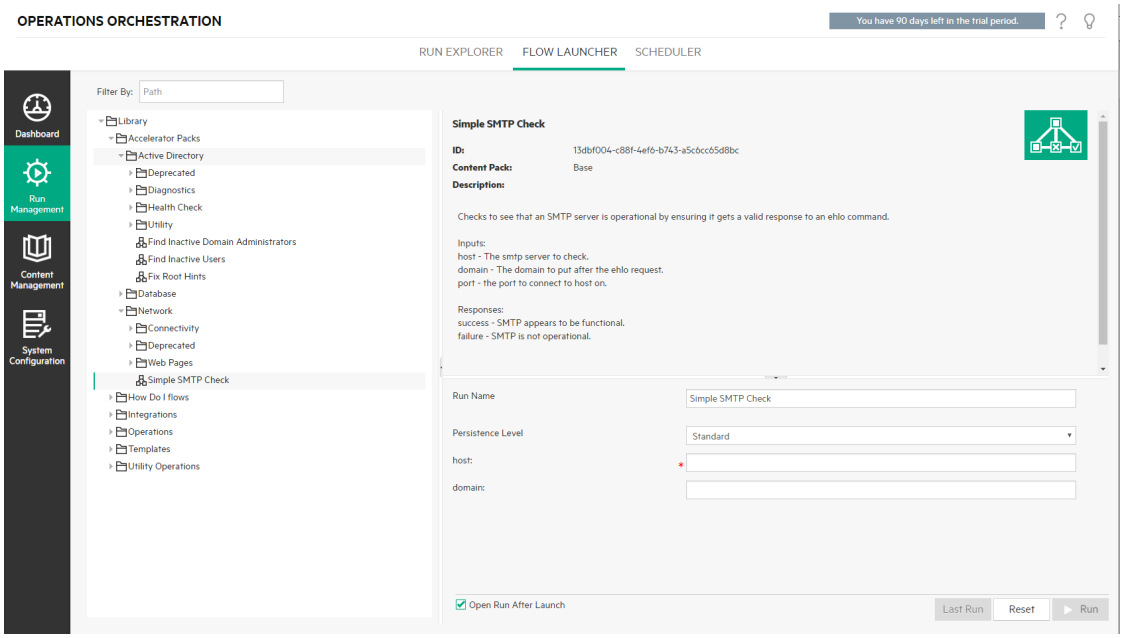

4. (Optional) To locate the flow that you need, enter part or all of the flow path in the **Filter By** text

box.

For example, if you enter "network", only flows that contain the word "network" or which are located in a path with a folder containing the word "network" are displayed.

> **Note:** The search is not case-sensitive.

5. Select the flow that you want to run. Details about the flow are displayed in the **Flow Details** pane on the right.
6. In the **Run Name** box, enter a name for the run.

   By default, the run name is the flow name, but you can change this. You may want to give it a name that identifies this particular run. For example, you may want to run the same flow multiple times, each with different input. Naming the runs makes it easier to identify them based on the inputs.
7. From the **Persistence Level** list, select the persistence level, in order to control the level of detail that is saved to the Run Log. The options are:
   - **Standard** - Includes raw results, step results, and big inputs/outputs.
   - **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs.

   > **Note:** The flow's default persistence level is selected by default.

   If you do not select a persistence level, the default is used. For information about how to set the default persistence level, see "Monitoring and Controlling Database Size" on page 112.
8. Enter any inputs that are required.

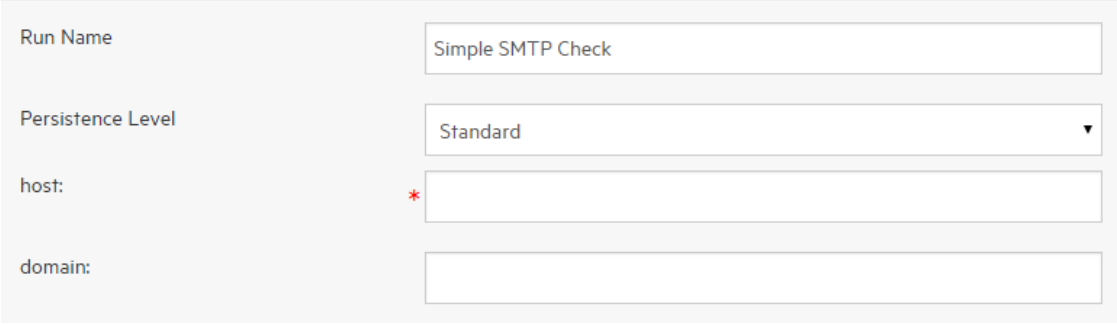

> **Note:** You can click the **Reset** button to apply the default values.

9. (Optional) Click the **Open Graph**  button to display a graphical representation of the flow.
10. (Optional) Select the **Open Run After Launch** check box to automatically track this flow after it is started.
11. Click **Run**.
    - If the **Open Run After Launch** check box was selected, the progress of the run is displayed in the Drill Down dialog box.

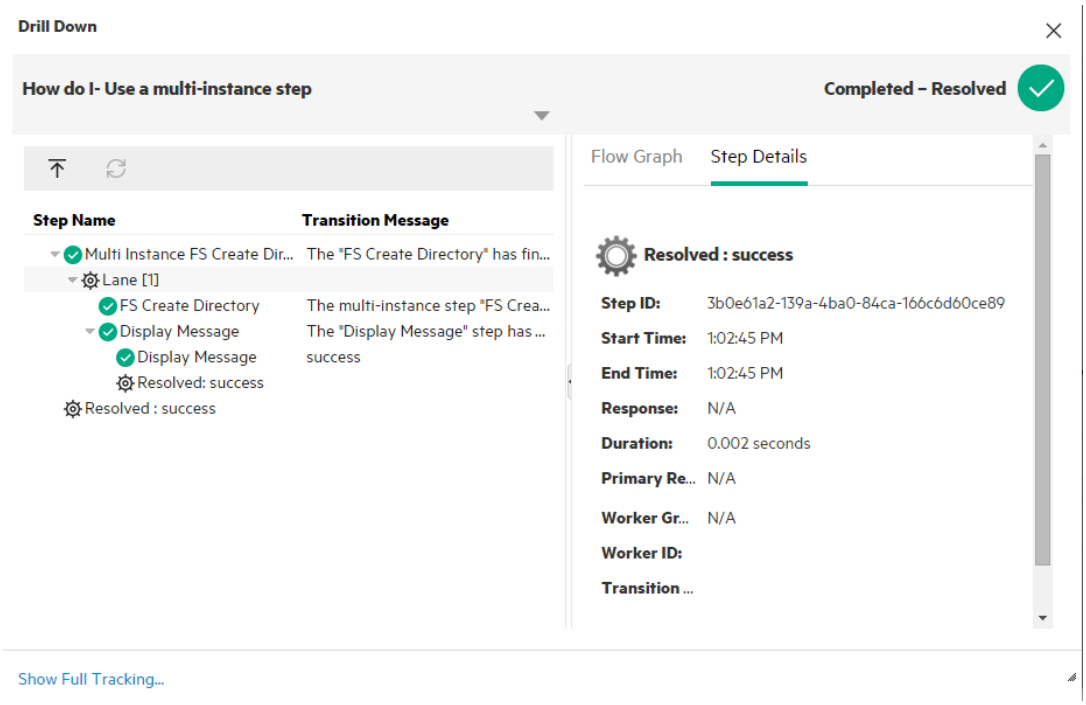

Click **Show Full Tracking** to navigate to the Run Explorer Drill Down. For more information about tracking the progress of a flow, see "Testing and Troubleshooting a Flow Run" on page 229.

- If the **Open Run After Launch** check box was not selected, the Drill Down dialog box is not opened. A message displays whether the flow was launched successfully.
- If you have run this flow previously, you can click the **Last Run** button to view the progress of your last run in the Drill Down dialog box.

Run a flow from the Flow Library

1. Click the **Content Management** button to display the Content Management workspace.
2. Click the **Flow Library** tab.
3. Browse for a flow in the flow library.

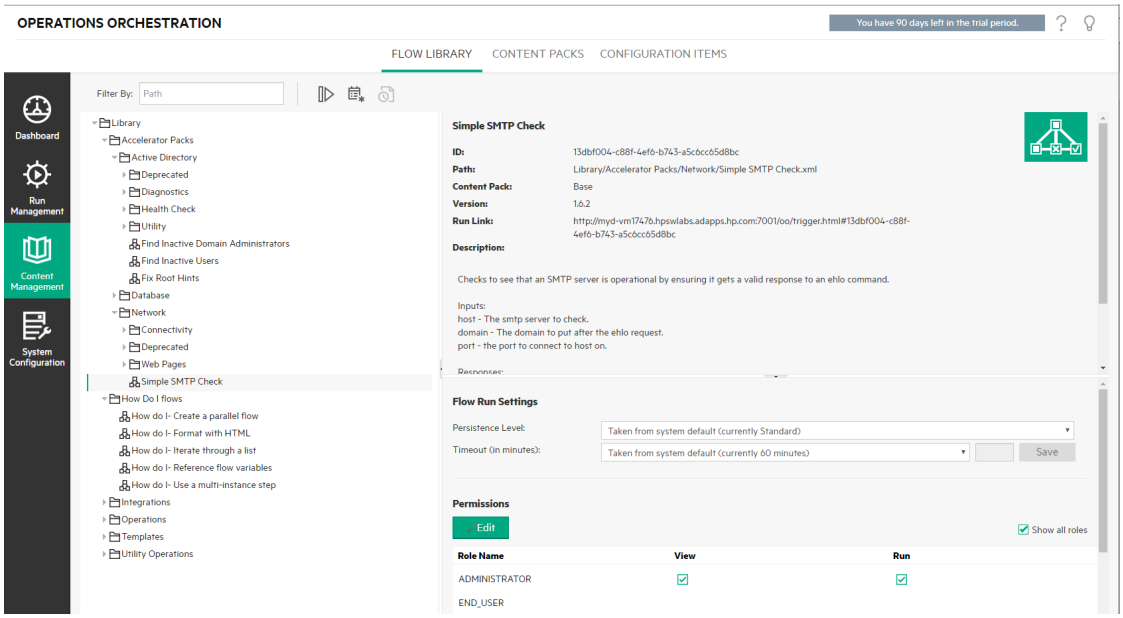

4. (Optional) To locate the flow that you need, enter part or all of the path in the **Filter By** text box.

   For example, if you enter "network", only flows that contain the word "network" or which are located in a path with a folder containing the word "network" are displayed.

   > **Note:** The search is not case-sensitive.

5. (Optional) Click the **Open Graph** button to display a graphical representation of the flow.

   

   Select the flow that you want to run, and click the **Run**  button in the toolbar.

6. In the Run Flow dialog box, enter a name for this flow run.

   By default, the run name is the flow name, but you can change this. You may want to give it a name that identifies this particular run. For example, you may want to run the same flow multiple

times, each with different input. Naming the runs makes it easier to identify them based on the inputs.

7. Enter the necessary inputs for the run in the Run Flow dialog box.

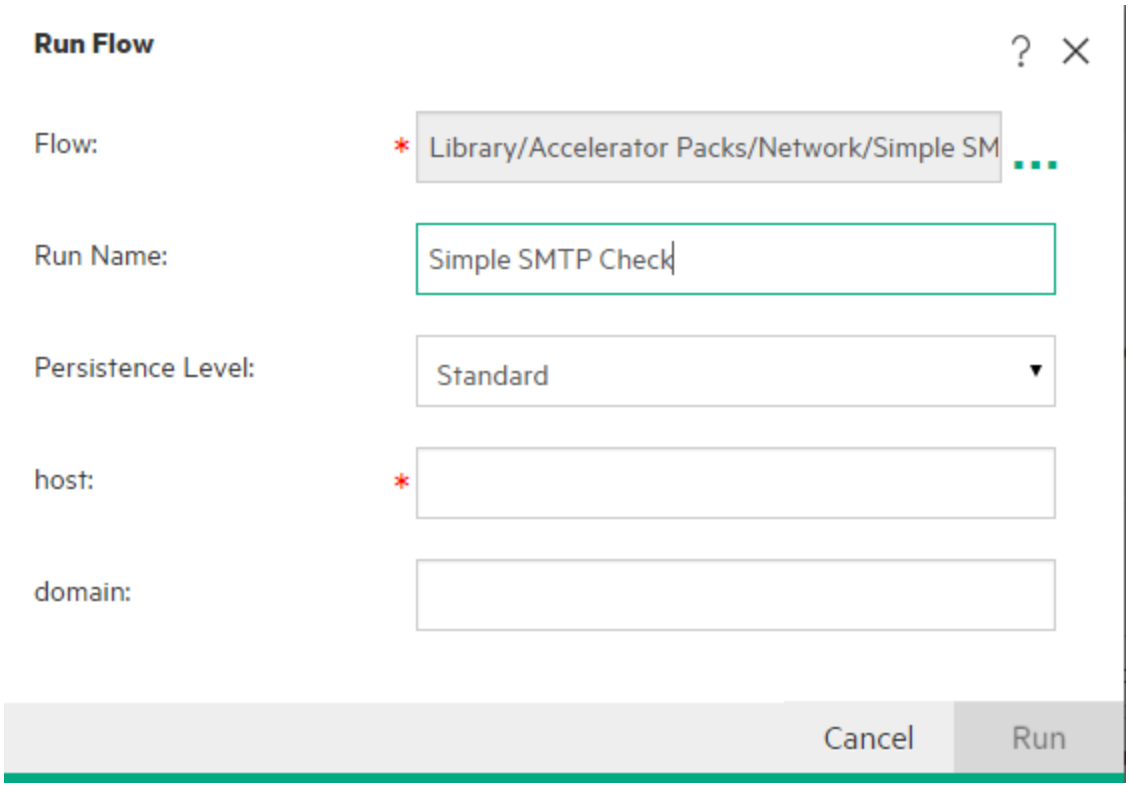

8. Click **Run**.

   The progress of the run is displayed in the Drill Down dialog box.

Run an interactive flow

1. Run the flow.
2. When the run reaches an interactive point (such as **Inputs Required**, **Handoff**, **Display Message**, and so on), its status is changed to **Paused - Action Required**.
3. The flow can be resumed from the Run Explorer Drill Down or from the Drill Down dialog box.
   - Select the run in the Run Explorer and click the **Drill down** button at the end of the row to display the Run Information window, and click the **Expand** button to expand the prompt message.
   - If the **Open Run After Launch** check box was selected, the prompt message appears in the

Drill Down dialog box.

> **Note:** You can click the **Contract** button on the left side of the Inputs Required dialog box to hide it again.

4. Read the message and follow its instructions to resume the run—enter the required inputs, handoff, and so on, and click **Resume**.

For information about the different statuses, see "Tracking and Managing Flow Runs" on page 215.

Run a flow with a gated transition

A gated transition restricts access to the next step to users who have been assigned a particular role.

> **Note:** Gated transition flows can only be resumed by users who have **Manage Others' Runs** permissions.

1. Run the flow.
2. When the flow run reaches the gated transition, if you have not been assigned a role with the required permission, the flow run is paused. The run status appears as **Paused – Action Required**.
3. If the **Open Run After Launch** check box was selected, the prompt message appears in the Drill Down dialog box.

   If not, select the run in the Run Explorer and click the **Drill down** button at the end of the row to display the Run Information window, and click the **Expand** button to expand the prompt message.
4. In the prompt message, click the **Hand off to another user (email this run)** link to create an email message to send to the user who needs to complete the flow run.
5. An email message opens with a message for the person to whom you are handing off the flow run and a link to the flow run. Enter the email address of this person and send the message.

Run a flow with a handoff transition

A handoff transition is one where a message appears, suggesting that the current user hand off the flow run to another user. Unlike with a gated transition, the first user has the option to resume the flow run without handing it off to another user.

1. Run the flow.
2. When the flow run reaches the handoff transition, the flow run is paused. The run status appears as **Paused – Action Required**.
3. If the **Open Run After Launch** check box was selected, the prompt message appears in the Drill Down dialog box.

   If not, click the **Drill down**  button at the end of the row to display the Run Information window, and click the **Expand**  button to expand the prompt message.
4. You have two options:
   - To hand off the flow run to another user, click **Hand off to another user (email this run)**. An email message opens with a message for the person to whom you are handing off the flow run and a link to the flow run. Enter the email address of this person and send the message.
   - Click **Resume** to continue the flow without handing it off.

Run a flow with a Paused - No Workers interaction

A flow has the **Paused – No workers** status when there is a worker group that has not yet been mapped to any workers. Before you can resume the run, you will need to map the worker group to actual workers. For more information see "Setting Up Configuration Items for a Content Pack" on page 164.

1. Run the flow.
2. When the flow run reaches the point where it relates to a worker group with no workers, the flow run is paused. The run status appears as **Paused – No Workers**.
3. If the **Open Run After Launch** check box was selected, the prompt message appears in the Drill Down dialog box.

   If not, click the **Drill down**  button at the end of the row to display the Run Information window, and click the **Expand**  button to expand the prompt message.
4. Map the worker group to the relevant workers. For details, see "Setting Up Configuration Items for a Content Pack" on page 164.
5. Return to the pause message and click **Resume**.

Reference Material

Flow Launcher tab

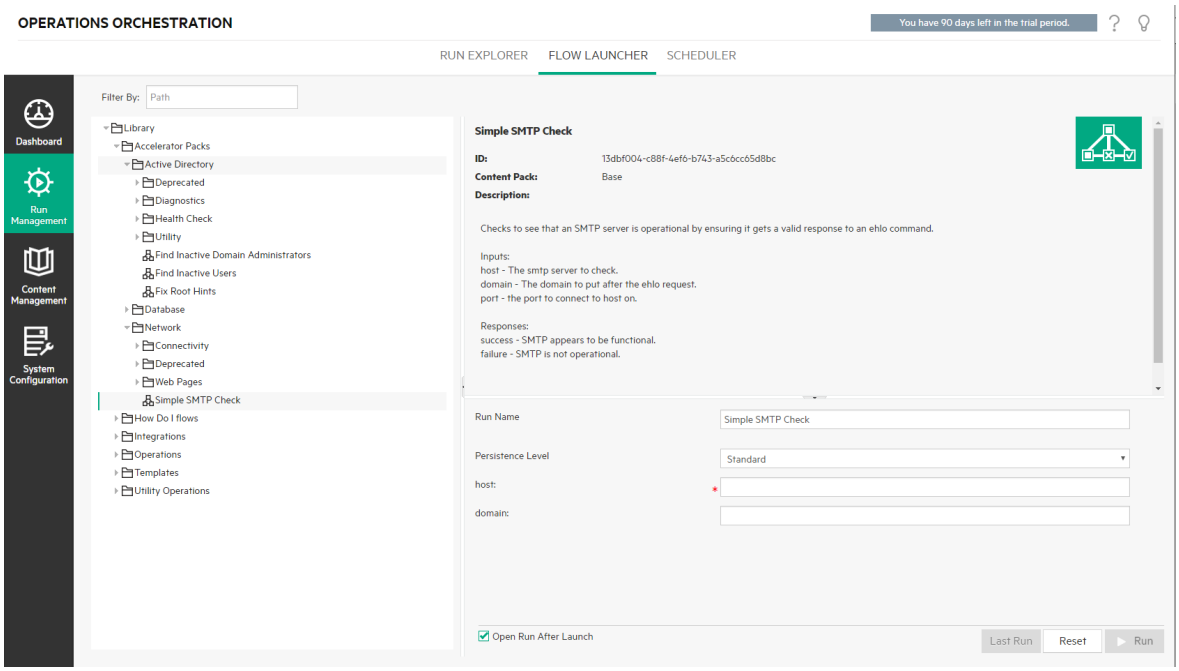

GUI item	Description
Filter By	To locate the flow that you need, enter part or all of the flow path in the **Filter By** text box to filter the flows.
Flows tree	Displays the flows that you have permission to view. Expand and collapse the nodes to see the flows in different folders.
Flow Details pane	Displays information about the selected flow.
Open Graph Button 	Click the **Open Graph** button to display a graphical representation of the flow. For more information, see "Viewing a Flow Graph" on page 159.
ID	Displays the UUID of the selected flow.
Content Pack	Displays the content pack that the selected flow came from.

Description	Displays a description of the selected flow. If required, use the scrollbar to scroll down and view the contents of the **Description** section. For OO Content flows, you can add input and output information under **Description**.
Input	Displays the input information for CloudSlang flows only. **Note:** This field may be encrypted as configured by the flow creator.
Output	Displays the output information for CloudSlang flows only. **Note:** This field may be encrypted as configured by the flow creator.
Inputs	Displays an explanation of the inputs required in the selected flow.
Run Name	In the **Run Name** box, enter a name for the run. The name does not have to be the same name as the original flow. You may want to give it a name that identifies this particular run. For example, you may want to run the same flow multiple times, each with different input. Naming the runs makes it easier to identify them based on the inputs.
Persistence Level	Select the persistence level, in order to control the level of detail that is saved to the Run Log. The options are: • **Standard** - Includes raw results, step results, and big inputs/outputs. • **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs. **Note:** The flow's default persistence level is selected by default. If you do not select a persistence level, the default is used. For information about how to set the default persistence level, see "Monitoring and Controlling Database Size" on page 112.
<Input fields>	Enter inputs if these are required. The number of input fields and the labels on them will vary, depending on the inputs that are required in the selected flow.
Open Run After Launch check box	Select to automatically track the flow after it is started. If this check box is selected, the flow's progress is displayed in the Run Explorer Drill Down dialog box.
Last Run button	If you have run this flow previously, you can click the **Last Run** button to view the progress of your last run in the Drill Down dialog box.
Reset button	Click the **Reset** button to clear all the inputs from the **Flow Input** pane.
Run button	Click the **Run** button to run the selected flow.

Run Flow dialog box

The Run Flow dialog box is displayed when the **Run** button is clicked, in the Flow Library.

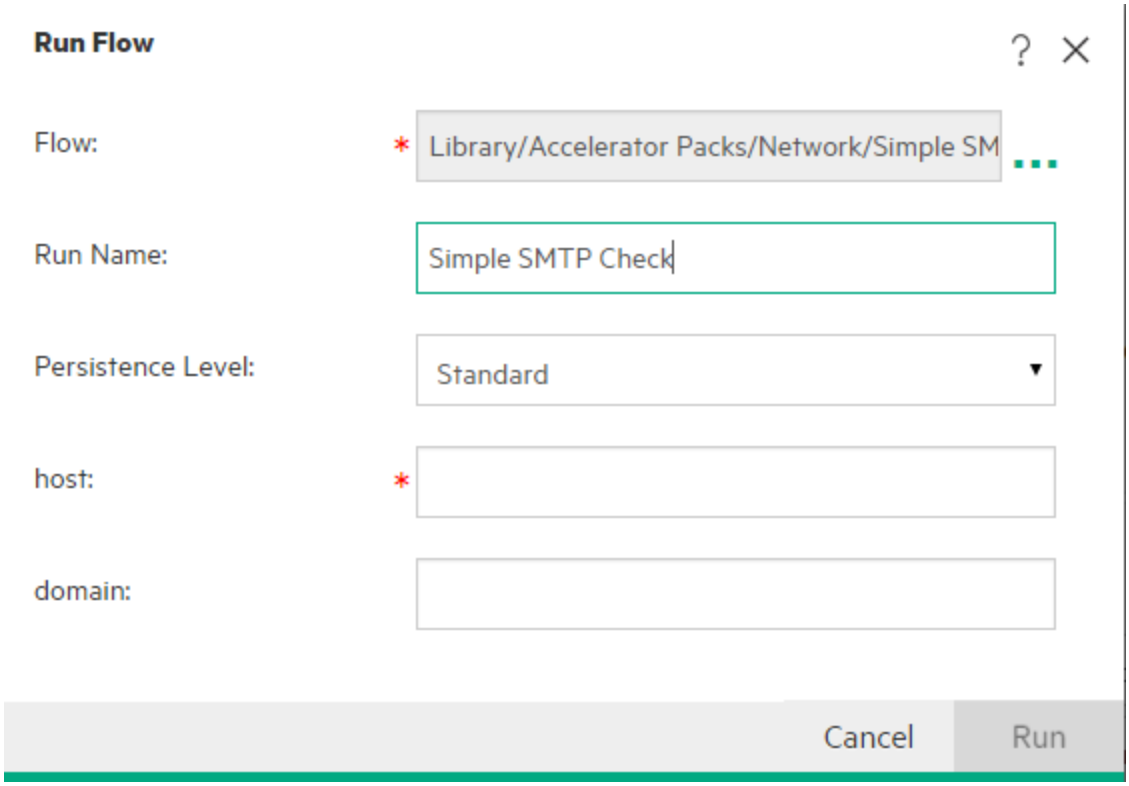

GUI item	Description
Flow	Displays the path to the flow.
Run Name	In the **Run Name** box, enter a name for the run. The name does not have to be the same name as the original flow. You may want to give it a name that identifies this particular run. For example, you may want to run the same flow multiple times, each with different input. Naming the runs makes it easier to identify them based on the inputs.
Persistence Level	Select the persistence level, in order to control the level of detail that is saved to the Run Log. The options are: • **Standard** - Includes raw results, step results, and big inputs/outputs. • **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs. For more information, see "Monitoring and Controlling Database Size" on page 112.

<Input fields>	Enter inputs if these are required. The number of input fields and the labels on them will vary, depending on the inputs that are required in the selected flow.
Run	Click to run the flow.

Drill Down dialog box

The Drill Down dialog box displays the progress of the run, when you run a flow from the Flow Launcher or from the Flow Library, when the **Open Run After Launch** check box is selected.

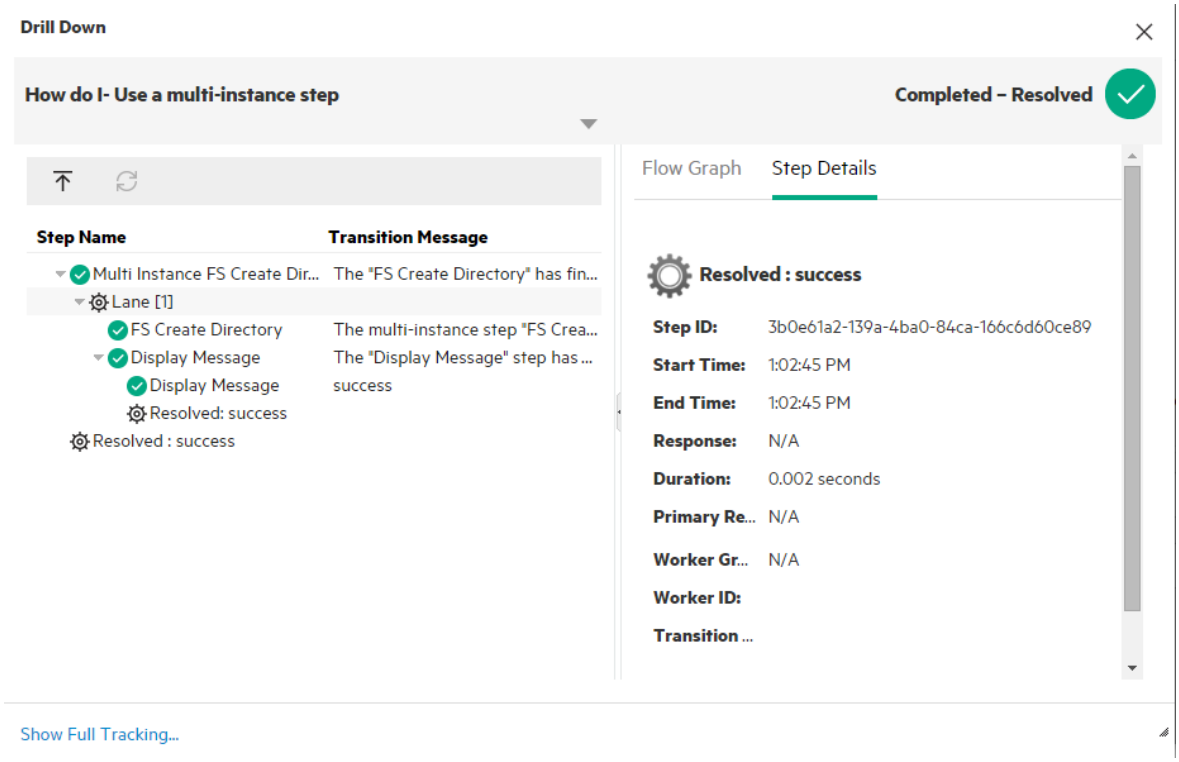

GUI item	Description
Show Full Tracking	Click **Show Full Tracking** to navigate to the Run Explorer Drill Down.
Other UI items	The UI items in the Drill Down dialog box are the same as those in the Run Explorer Drill Down. For more information about the Run Explorer Drill Down, see "Testing and Troubleshooting a Flow Run" on page 229.

Rerunning a Flow from a Specific Step

Central provides the ability to rerun a flow from a specific step. This is useful when recovering complex flows, where you do not want to rerun the entire flow from the beginning if it failed or was canceled.

The administrator enables/disables the rerun functionality using the **Enable the rerun option for failed/canceled runs** check box in **System Configuration > System Settings**.

You can rerun the flow from a specific step by clicking on the rerun icon . For example:

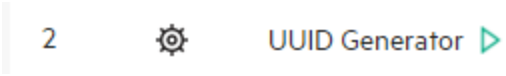

In order to rerun a flow from a specific step, you must first mark the step as a rerun point in Studio. See "Enable Rerunning a Flow from a Specific Step" in the *HPE OO Studio Authoring Guide* .

Note:

- You cannot rerun steps inside parallel or multi-instance lanes. You will not see the rerun icon for these types of steps.
- You cannot rerun steps in flows executed via a remote debugger. You will not see the rerun icon for these flows.
- This functionality is available in HPE OO version 10.70 and above. If you use the new Rerun functionality in a flow, you cannot deploy the flow on earlier Central versions.
- Defining rerun points means that more data is persisted in the database.
- Marking a large number of steps as rerun points increases OO resource consumption and may impact the overall flow execution performance. If you want to define many rerun points, it is recommended that you increase the maximum Java Heap Size. See "Increasing the JVM Heap Size" in the *HPE OO Tuning Guide* .

What do you want to do?

Enable the rerun option

1. Select the **System Configuration > System Settings** tab.
2. In the Run Settings area, select the **Enable the rerun option for failed or canceled runs** check

box.

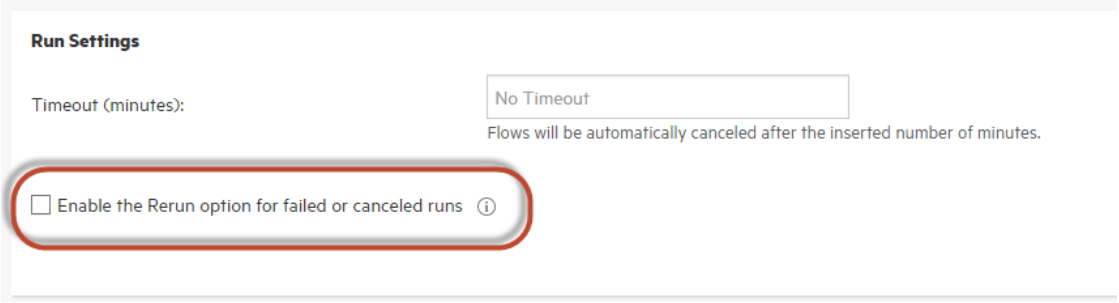

3. Reselect the check box to disable the rerun functionality. The rerun icon will no longer appear next to steps that can be rerun.

Rerun a flow from a specific step

1. In Studio, mark the step as a rerun point in Studio. See "Enable Rerunning a Flow from a Specific Step" in the *HPE OO Studio Authoring Guide* .
2. In the **Run Management** workspace, select the desired flow.
3. Double click on the flow or click the **Drill down** button at the end of the row to drill down into the flow steps.
4. Click on the step from which you want to start the flow rerun.
5. Click the rerun icon next to the step.

   The Rerun window opens:

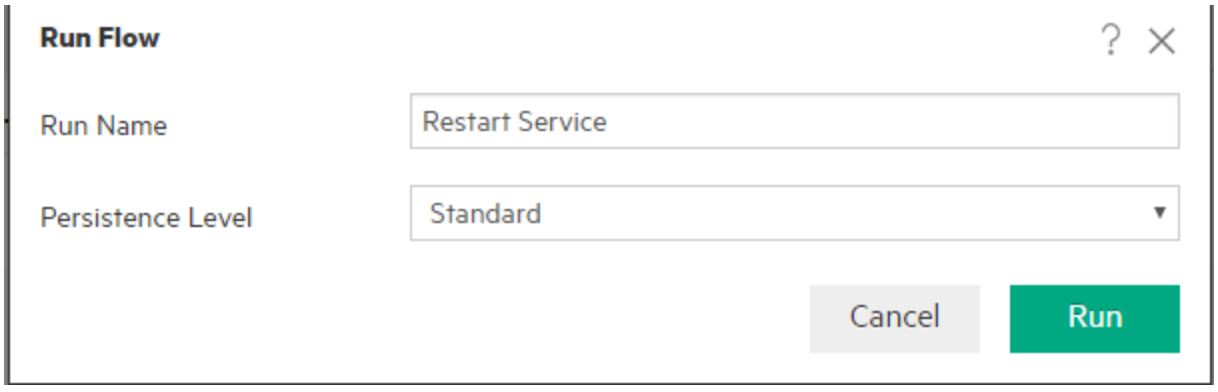

6. Select the Persistence Level from the dropdown list and click **Run**.

For details on how to define the persistence levels, see "Default Persistence Level for the Run Log " on page 113.

Drill down into a Rerun Step

1. Click on the arrow next to the rerun step to drill down into the details of the rerun step.

   For example:

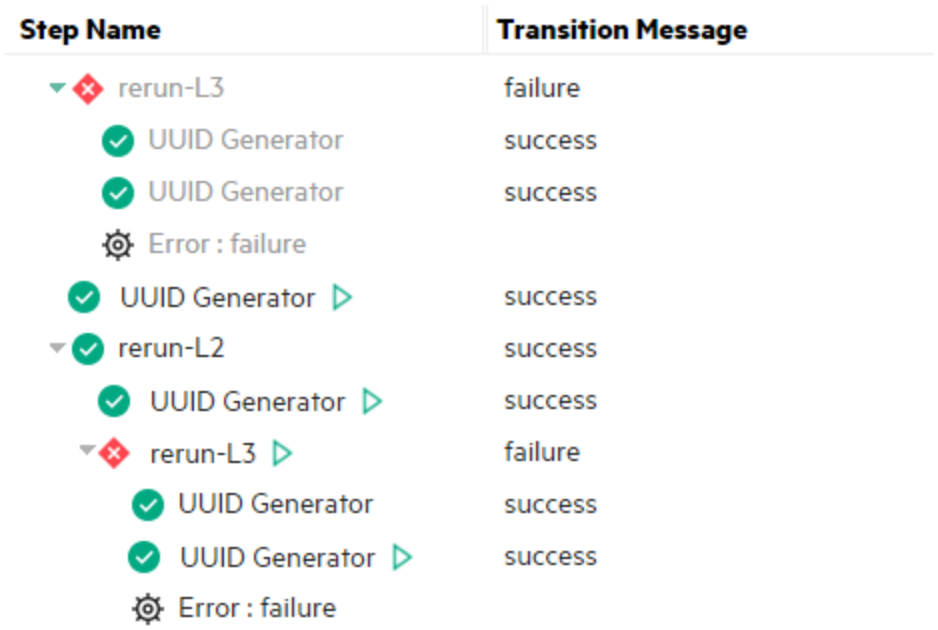

   The substeps of the rerun step are shown.

2. Click Step Details in the right pane to see information on the selected substep.

Reference Material

For reference information about the **Run Explorer** tab, see "Tracking and Managing Flow Runs" on page 215.

Scheduling Flow Runs

Schedules allow you to control when to run your flows. You can specify a schedule to run for a specific flow. You can also set up recurring schedules for a task that needs to be repeated.

For example, you need to regularly check whether a number of servers are online. In this case, you can define a flow to check the IP address and then create a schedule to run that flow.

Schedules are displayed in the **Scheduler** tab. The display of schedules depends on the permissions that you have been assigned in your role:

- If you have been assigned a role with the **View Schedules** permission, you will be able to view the HPE OO schedules and their details.
- If you have been assigned a role with the **Manage Schedules** permission, you will be able to view and edit schedules.

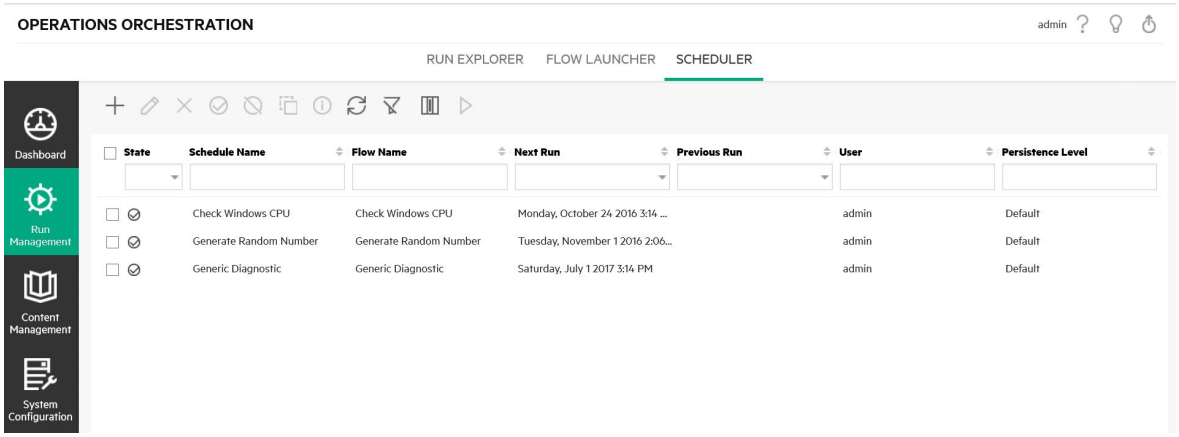

> **Note:** The Central Scheduler does not work when SAML authentication is used.

Schedule Owner

Each schedule has an owner. This is the user who created the schedule. The name of the owner appears in the **User** column in the Scheduler

- In the **Scheduler** tab, the **User** column indicates the schedule owner.
- In the **Run Explorer** tab, the **User** column indicates the owner of this flow execution (the user who performs the scheduled run).

If another user updates the schedule (for example, the recurrence, name, or time zone), the schedule owner changes to this user. This user is displayed as the schedule owner in the Scheduler and as the flow executor in the Run Explorer.

> **Note:** The new owner appears only for runs that started after the change. For runs that started before the change, the original owner is displayed in the **User** column.

If the schedule is created or updated by a user who is logged in as an anonymous user, the **User** column will display N/A.

> **Note:** If the schedule is updated by a user without run permissions on the scheduled flow, this flow will not run.

> **Note:** If an LDAP configuration exists and the configured LDAP server is not responding, this will affect all types of schedules. This includes both schedules that are run with a LDAP user and schedules that are run with an internal user.

Schedule Name

Each schedule has a name. The flow name is entered by default, but you may prefer to use a different name, for example, if you have multiple schedules using the same flow.

> **Note:** If you give the schedule a different name, this is the name that will appear in the Run Explorer, rather than the flow name.

Scheduler Troubleshooting

A scheduler that fails to trigger is known as a misfire.

You can set the action (instruction) to take place upon a scheduler misfire in the file **<OO installation folder>\central\tomcat\webapps\oo\WEB-INF\classes\META-INF\spring\quartzContext.xml**.

In this file, set the **p:misfireInstructionName** attribute for simple trigger and cronTrigger beans. The default instruction is "MISFIRE_INSTRUCTION_DO_NOTHING"

In the case of a misfire, the following log message opens:

```
There was a misfire in the scheduler <name_of_the_scheduler>
```

> **Note:** The logging of a scheduler misfire will occur only when the Central service is running.

What do you want to do?

View existing schedules

To view the **Scheduler** tab, you must be logged in with a user account that is assigned a role with **View schedules** permissions.

1. Click the **Run Management** button to display the Run Management workspace.
2. Click the **Scheduler** tab. A table displays all the current schedules, including the following information:
   - **State**: Displays whether the schedule is enabled or disabled .
   - **Schedule Name**: Displays the name that was given to this schedule. This value can be edited in the Edit Schedule dialog box.
   - **Flow Name**: Displays the name of the flow that the schedule relates to. This value is taken from the flow and cannot be modified in the schedule.
   - **Next Run**: Displays the time and date when the next run is scheduled to occur. The time and date displayed is in the your time zone.
   - **Previous Run**: Displays the time and date when the previous run occurred, if this exists.
   - **User**: Displays the name of the user who created or modified this schedule.

Display the details of a schedule

You can drill down to see the details of a schedule in the Schedule Details dialog box. This functionality is useful for users with **View Schedules** permission, but without permission to edit schedules.

The details include information that is not displayed in the table in the **Scheduler** tab, such as recurrence pattern, range of recurrence, and the timezone in which the schedule runs.

To display the details of a schedule:

1. Click the **Scheduler** tab.
2. Select the check box next to the schedule that you want to view and click the **Details** button in the **Scheduler** toolbar. The details of the selected schedule are displayed in the Schedule Details dialog box.

> **Note:** The **Details** button is only available if a single schedule is selected.

3. When you have finished viewing the details of the schedule, click **OK**.

Edit schedules

To edit a schedule, you must be logged in with a user account that is assigned a role with **Manage schedules** permission.

To edit a schedule:

1. Click the **Scheduler** tab.
2. Select the check box next to the schedule that you want to edit and click the **Edit**  button in the **Scheduler** toolbar.
3. In the Edit Schedule dialog box, make the required changes to the schedule, and then click **Save**.

> **Note:** If you try to set the end date for the schedule to the current date (or earlier), the Save button is disabled and the following warning message appears:
>
> `The end date must be after the start date.`

Clone schedules

You can create an identical copy of a schedule that appears in the **Scheduler** tab. You can make any required changes to this schedule in the Clone Schedule dialog box.

To clone a schedule:

1. Click the **Scheduler** tab.
2. Select the check box next to the schedule that you want to clone, and click the **Clone**  button in the **Scheduler** toolbar.

   > **Note:** The **Clone**  button is only enabled if a single schedule is selected.

   The Clone Schedule dialog box opens. All the fields in the dialog box are populated with the values from the selected schedule.
3. If required, edit the schedule details in the Clone Schedule dialog box. You may wish to give the clone a different name.
4. Click **Save**.

> **Note:** If you try to set the end date for the cloned schedule to the current date (or earlier), the Save button is disabled and the following warning message appears:
>
> `The end date must be after the start date.`

Delete schedules

You can delete individual or multiple schedules.

To delete schedules:

1. Click the **Scheduler** tab.
2. Select the check boxes next to the schedules that you want to delete and click the **Delete**  button in the **Scheduler** toolbar.
3. In the confirmation dialog box, click **Yes**.

The selected schedules are deleted from the **Scheduler** tab.

Enable a schedule

By default, when a new schedule is created, it is enabled. If a schedule has been disabled, you need to enable it before it will run. You can enable individual or multiple schedules.

To enable a schedule for a flow run:

1. In the **Scheduler** tab, select the check box next to each disabled schedule that you want to enable.
2. In the **Scheduler** toolbar, click **Enable**  .

   > **Note:** The **Enable** button is only available if at least one disabled schedule is selected.
3. In the confirmation dialog box, click **Yes**.

The icon in the **State** column changes to **Enabled** for this schedule.

Disable a schedule

You can disable individual or multiple schedules. These schedules remain in the **Scheduler** tab, but they will not run until they are enabled.

To disable a schedule for a flow run:

1. In the **Scheduler** tab, select the check box next to each enabled schedule that you want to disable.

2. In the **Scheduler** toolbar, click **Disable**  .

> **Note:** The **Disable** button is only available if at least one enabled schedule is selected.

3. In the confirmation dialog box, click **Yes**.

The icon in the **State** column changes to **Disabled** for this schedule.

Filter the display of schedules

You can filter the schedules that are displayed in the **Scheduler** tab. This can help you locate schedules quickly, when there is a large number of schedules.

For example, if you are looking for schedules of flows relating to connectivity, you could enter "Connect" as the filter term in the **Flow Name**column.

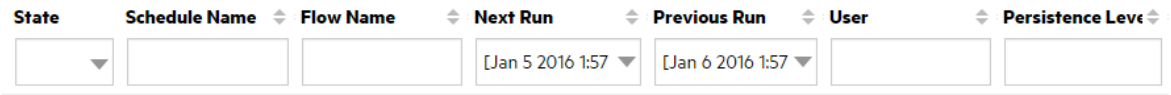

> **Note:** The filter is not case-sensitive.

1. Click the **Scheduler** tab.
2. In the text boxes under the column titles, enter the information to filter by.

   If you do not know the exact name, enter part of the name.

Refresh the display of schedules

If someone creates, edits, deletes, enables, or disables a schedule in another client, this does not show in the **Scheduler** tab until you have refreshed the page.

To refresh the display of schedules:

1. Click the **Scheduler** tab.
2. In the **Scheduler** toolbar, click **Refresh**  .

> **Note:** It is also possible to refresh the page using a browser refresh command (for example, using key combinations such as F5 or Ctrl+R), but this loads the entire page. Clicking the **Refresh** button is faster, because it only refreshes the display of schedule data .

Sort the display of schedules

If there are multiple schedules displayed in the **Scheduler** tab, you can sort them according to any of the column headings.

1. Click the **Scheduler** tab.
2. Click the column header that you want to use to sort the schedules.

Run a Scheduled Flow Now

You can trigger an ad-hoc run of a scheduled flow.

Central users that can perform this action must have the following permissions enabled for their role:

- View Schedules and Manage Schedules permissions. For more information see "Setting Up Security – Roles" on page 32.
- View and Run content permissions. For more information, see "Setting Permissions for Content" on page 150.

> **Note:** If a schedule is set to end after a specific number of occurrences, running this scheduled flow ad-hoc will not be counted as an occurrence. Only runs that are initiated by the schedule's recurrence pattern will be counted in the number of occurrences.

1. Click the **Scheduler** tab.
2. Select the schedule from the list.
3. In the Scheduler toolbar, click the Run-now button  .

Create a schedule

To create a schedule, you must be logged in with a user account that is assigned a role with **View schedules** and **Manage schedules** permissions.

1. You can create a new schedule either from the Scheduler or from the Flow Library:
   - To create a schedule from the Flow Library:
     i. Click the **Content Management** button to display the Content Management workspace.
     ii. Click the **Flow Library** tab.
     iii. Select a flow from the Flow Library.
     iv. Click the **Schedule**  button. The Create Schedule dialog box opens.

v. Click the browse ▪▪▪ button to browse for and select the flow for which you want to create the schedule.

- To create a schedule from the Scheduler:

i. Click the **Run Management** button to display the Run Management workspace.

ii. Click the **Scheduler** tab.

iii. In the **Scheduler** toolbar, click the **Create** + button.

iv. Select the flow for which you want to create the schedule, and click **OK**. The Create Schedule dialog box opens.

2. In the **Schedule Name** box, enter a meaningful name for this schedule. The flow name is entered by default, but you may prefer to use a different name, for example, if you have multiple schedules using the same flow.

> **Note:** If you give the schedule a different name, this is the name that will appear in the Run Explorer, rather than the flow name.

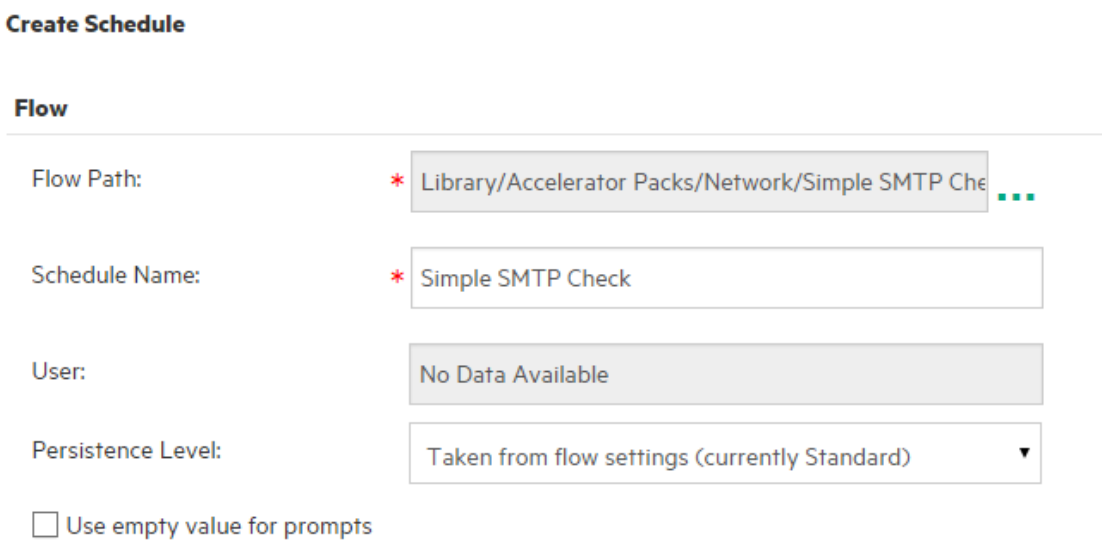

4. From the **Persistence Level** list, select the persistence level for the scheduled run, to control the level of detail that is saved to the Run Log. The options are:

- **Standard** - Includes raw results, step results, and big inputs/outputs.
- **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs.

- **Taken from flow settings** - whatever was selected as the default persistence level for the flow.

3. Select the **Use empty value for prompts** check box if you want the schedule to skip over prompts for inputs. For any prompt that has no value assigned, it will use blank, and the scheduled flow will run without stopping to prompt for values.

   However, note that if any of these are required inputs, the flow will fail.

   This option is selected by default after an upgrade from version 9.x to 10.x. This is to enable backward compatibility, because this is how all schedules worked in 9.x.

   If you do not select this check box, the run is stopped and prompts for inputs, even if they are not required inputs.

4. In the **Recurrence pattern** section, select from one of the following recurrence frequencies for the schedule: daily, weekly, monthly.

   - **Daily**: To run the flow on a daily basis, select **Daily** and one of the following:

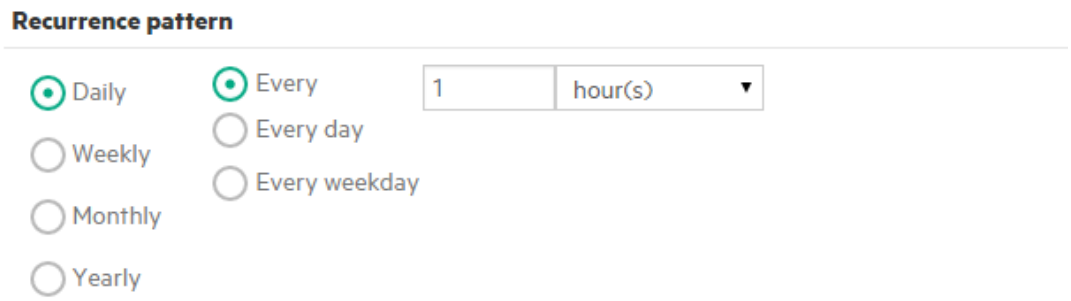

   - **Every**: To run the flow every day at regular intervals, select **Every**, and then enter the value and select **minutes** or **hours**.

     > **Note:** If you select **Every**, the interval field cannot be left empty, and the value must be at least **1**.

   - **Every day**: Select to run the flow once every day of the week.

     > **Note:** The flow run will start at the time set in the **Start time** field in the **Range of recurrence** section.

   - **Every weekday**: Select to run the flow once a day, on weekdays only.

     > **Note:** If the working week is not Monday-to-Friday, use the **Weekly** option instead, and select the working days.

- **Weekly**: To run the flow on a weekly basis, select **Weekly** and then select the days of the week on which you want the flow to run.

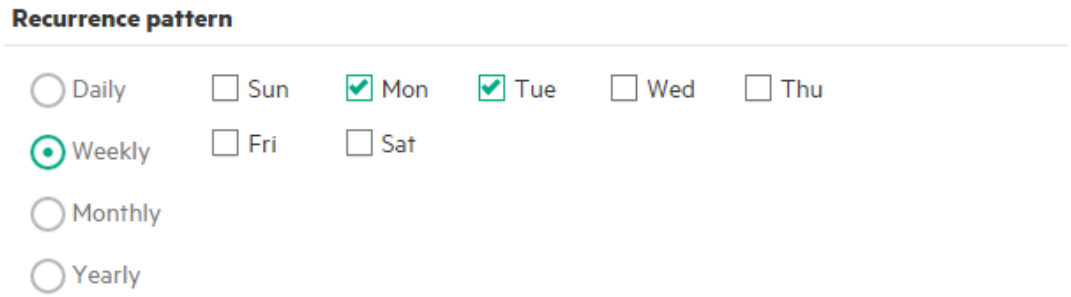

> **Note:** If you select **Weekly**, you must select at least one day.

- **Monthly**: To run the flow every month, select **Monthly** and one of the following:

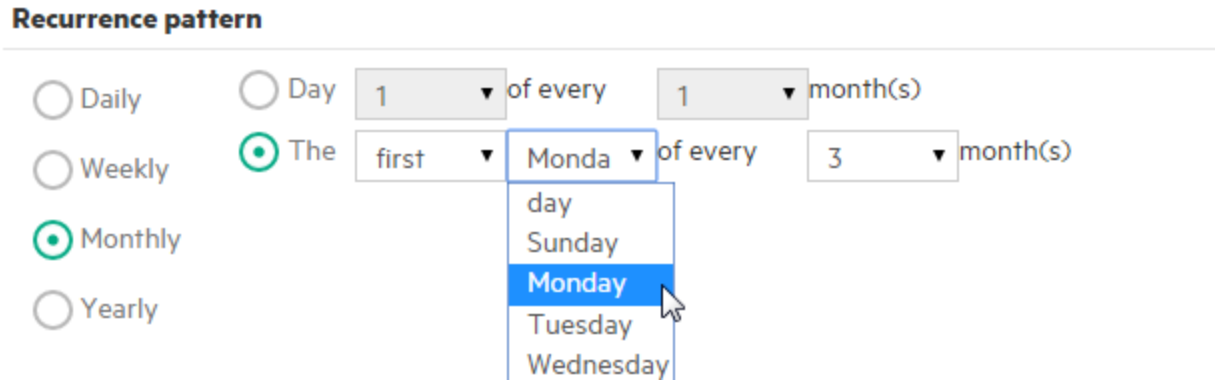

  - To run the flow on a particular date in each month, select **Day** and select the date. For example, select **15** to run the flow on the 15th of every month.
  - To run the flow once a month on a particular day of the week, select the day of the week and whether it is the first, second, third, fourth, or last time that this day occurs in the month. For example, to run the schedule on the third Friday of each month, select **third**, then select **Friday**, and then select **1** month(s).
  - To run the flow at intervals of several months, select a number of months. For example, to run a schedule once a quarter, on the first Monday of the quarter, select **first**, then select **Monday**, and then select **3** months.

    For example, in a schedule starting on April 1, 2013, and ending on December 20, 2013, with a recurrence pattern of the third day of every 2 months, the flow will run on June 3, August 3, October 3, and December 3.

- To run the flow on the last day of the month, select **last** and then select **day** instead of a specific day of the week.

◦ **Yearly**: To run the flow once a year, select **Yearly** and one of the following:

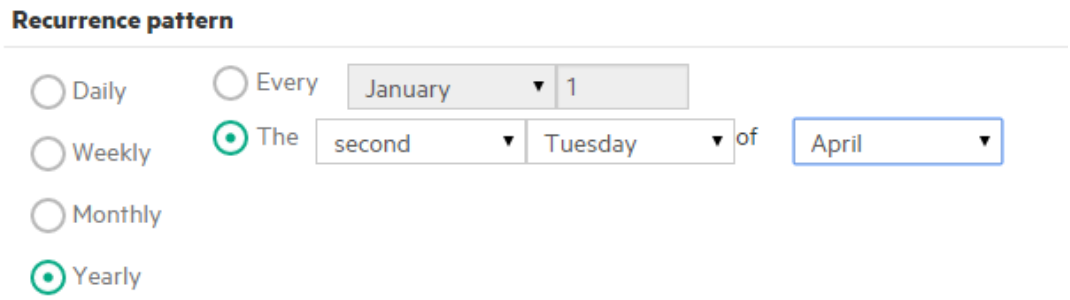

- To run the flow on a particular date each year, select **Every** and then select the month and the date. For example, every July 4.
- To run the flow once a year on a particular month and on a particular day of the week, select the month, then select the day of the week and whether it is the first, second, third, fourth, or last time that this day occurs in the month. For example, the second Monday of each November.

5. In the **Range of recurrence** section, in the **Start date** box, enter the date on which you want the flow schedule to start.

> **Note:** If a schedule with an earlier date undergoes content upgrade, the start date is changed to the current date.

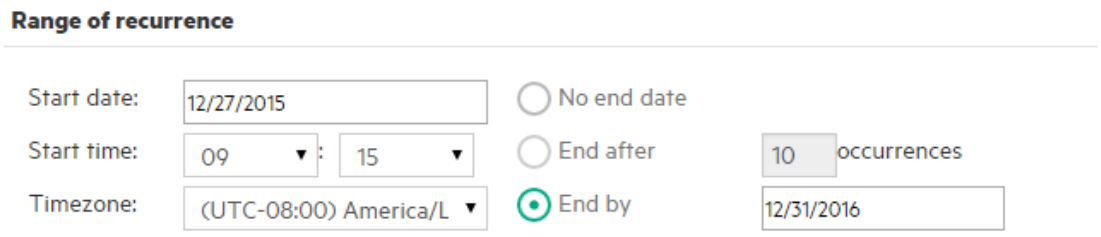

> **Note:** The date format appears in accordance with the locale of the client browser.

6. In the **Start time** boxes, enter the time at which you want the flow schedule to start, in hours and minutes.
7. From the **Timezone** list, select the time zone where you are located.

> **Note:** Time zones around the world are expressed as positive or negative offsets from

> Coordinated Universal Time (UTC). When you create a schedule, a time zone is selected by default from the alphabetically-ordered list that is currently available. HPE OO selects the first time zone from the list that has the offset equal to the client offset (computed with daylight savings, if relevant).
>
> Note that this timezone might not be geographically correct. For example, if you are located in a location where the time zone is UTC-09.00, the default time zone will be **America/Anchorage**. If desired, you can select the geographically correct location from the **Timezone** list.

8. Select when you want the schedule to end:
   - Select **No end date** if you do not want to set an end date for the schedule.
   - To limit the schedule to a specific number of runs, select **End after** and enter the number of runs after which the schedule will stop. This option is only available for daily schedules that have been set to run at intervals of a number of hours or minutes.
   - To end the schedule on a specific date, select **End by** and enter the date after which schedule will stop.
   - > **Note:** If you try to set the end date for the schedule to the current date (or earlier), the Save button is disabled and the following warning message appears:
     >
     > `The end date must be after the start date.`
9. If the flow requires inputs, enter them in the **Inputs** section on the right side. If a red star appears next to an input box, this means that the input is required, and the **Save** button will not be enabled unless an input value is provided.
10. If you want to enter additional inputs that are not flow inputs, click **Add Input** at the bottom of the **Inputs** section, and enter the name and value of the input.

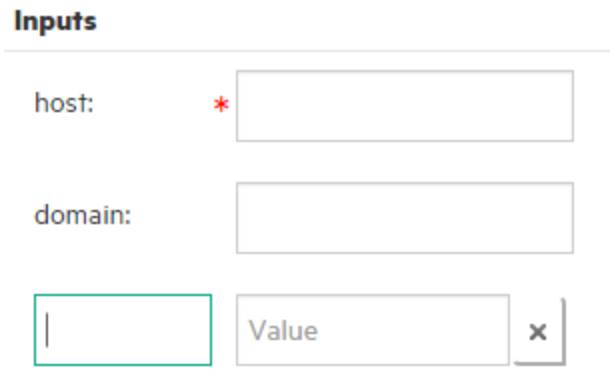

11. Click **Save**.

> **Note:** If the **Save** button is not enabled, this means that a required input has not been provided, or there is a mistake in one of the fields in the dialog box. For example, the **Save** button is not enabled if a non-existent date is entered, such as February 31, or if no days of the week are selected for a weekly recurrence pattern.

Reference Material

Scheduler tab

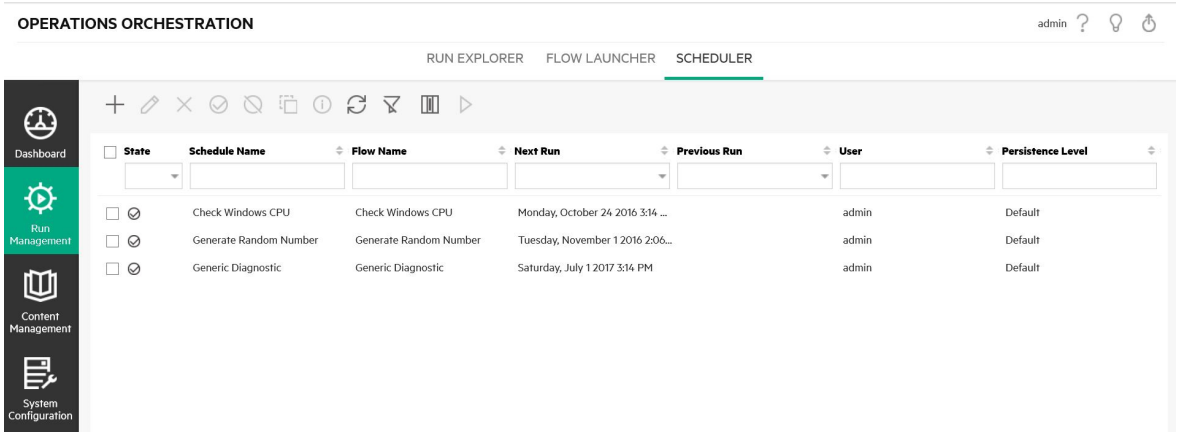

GUI item	Description
New button	Opens the Create Schedule dialog box, so that you can add a new schedule.
Edit button	Opens the Edit Schedule dialog box, so that you can edit the selected schedule. The **Edit** button is only available if at least one schedule is selected.
Delete button	Deletes the selected schedules. The **Delete** button is only available if at least one schedule is selected.
Enable button	Enables the selected schedules. The **Enable** button is only available if at least one disabled schedule is selected.
Disable button	Disables the selected schedules. The **Disable** button is only available if at least one enabled schedule is selected.

Clone button 	Creates an identical copy of the selected schedule. The **Clone** button is only enabled if a single schedule is selected.
Details button 	Displays the details of the selected schedule. The **Details** button is only enabled if a single schedule is selected.
Refresh button 	Click to refresh the display of schedules.
Select Columns button 	Click to display the column picker, to choose which columns to display.
Clear Filters button 	Click to remove the filters and display all the schedules.
Run now button 	Click to trigger an ad-hoc run of the scheduled flow. **Note:** The "Run Now" button is only enabled if a single schedule is selected.
Filter boxes	In the filter boxes at the top of the columns, enter text to filter the displayed schedules.
State	Displays whether the schedule is enabled  or disabled  .
Schedule Name	Displays the name that was given to the schedule. In the filter box under **Schedule Name**, enter the name of a schedule or part of a name, to filter the displayed schedules.
Flow Name	Displays the name of the flow that is run in the schedule. In the filter box under **Flow Name**, enter the name of a flow or part of a name, to filter the displayed schedules.
Next Run	Displays the time when the next run is scheduled.
Previous Run	Displays the time when the previous run in the schedule occurred.
Persistence Level	Displays the persistence level, which controls the level of detail that is saved to the Run Log: • **Standard** - Includes raw results, step results, and big inputs/outputs. • **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs.

	• **Taken from flow settings** - whatever was selected as the default persistence level for the flow. For more information, see "Monitoring and Controlling Database Size" on page 112.
User	Displays the name of the user who created or modified the schedule. If the schedule is created or updated by a user who is logged in as an anonymous user, the **User** column will display N/A.

Create Schedule dialog box

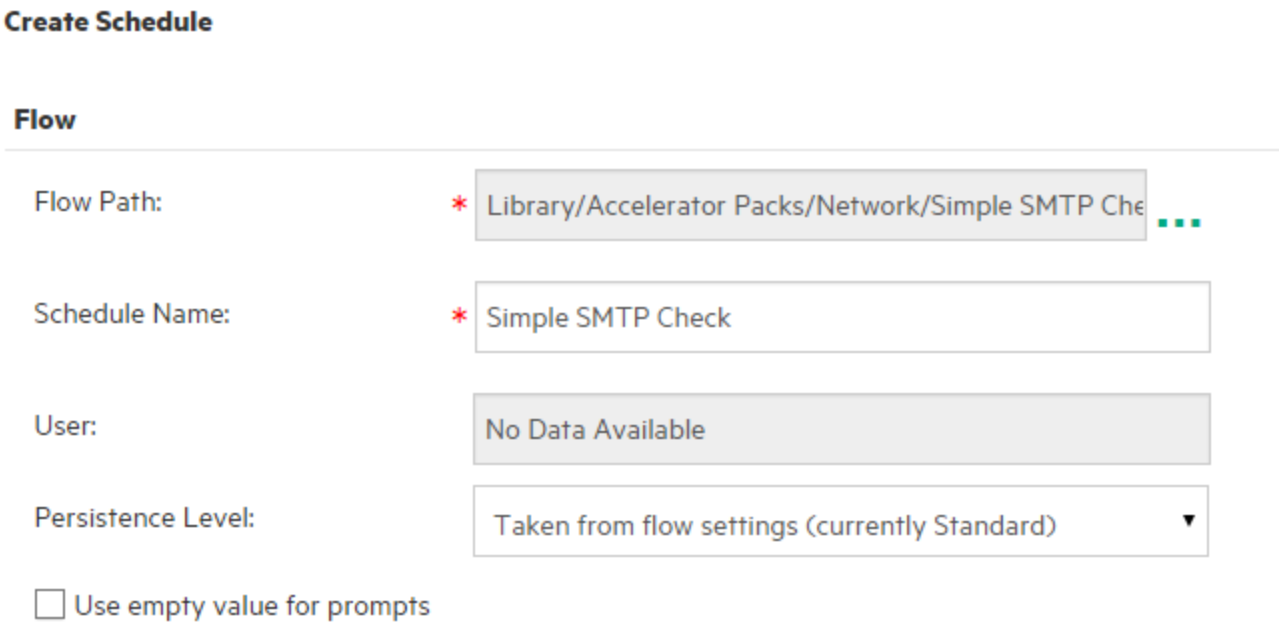

GUI item	Description
Flow Path	Displays the path to where the flow is located.
Schedule Name	Enter a meaningful name for this schedule.
Use empty value for prompts	Select this check box if you want the schedule to skip over prompts for inputs. For example, when you just want the schedule to run without you, and you don't want to wait to see the prompt and enter the input. However, note that if input is required for that flow, the flow will fail. This option is selected by default after an upgrade from version 9.x to 10.x. This is to

	enable backward compatibility, because this is how all schedules worked in 9.x
Persistence Level	Select the persistence level, in order to control the level of detail that is saved to the Run Log. The options are: • **Standard** - Includes raw results, step results, and big inputs/outputs. • **Extended** - Includes the complete raw results, step results, subflow outputs (from step results), and big inputs/outputs. • **Taken from flow settings** - whatever was selected as the default persistence level for the flow. For more information, see "Monitoring and Controlling Database Size" on page 112.
Recurrence pattern section	In the **Recurrence pattern** section, you define the frequency at which the schedule recurs. The **Recurrence pattern** section changes, depending on whether you select **Daily**, **Weekly**, **Monthly**, or **Yearly**.
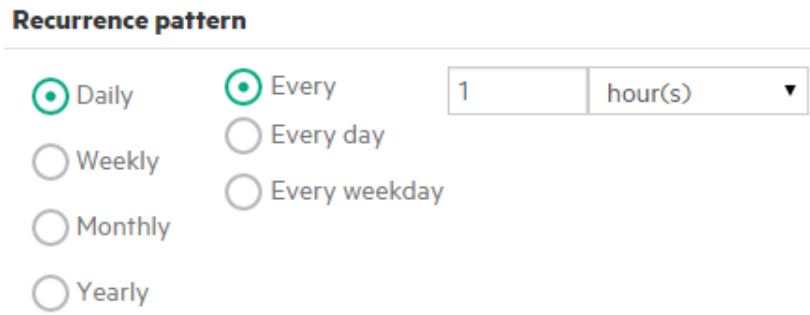	
Daily - Every <n> minutes/hours	To run the flow every day at regular intervals, select **Daily**, select **Every**, enter the value, and select **minutes** or **hours**. **Note:** If you select **Every**, the interval field cannot be left empty, and the value must be at least **1**.
Daily – Every day	To run the flow once a day, every day of the

	week, select **Daily**, and then select **Every day**. **Note:** The flow run will start at the time set in the **Start time** field in the **Range of recurrence** section.
Daily – Every weekday	To run the flow every weekday, select **Daily**, and then select **Every day**. **Note:** If the working week is not Monday-to-Friday, use the **Weekly** option instead, and select the working days.
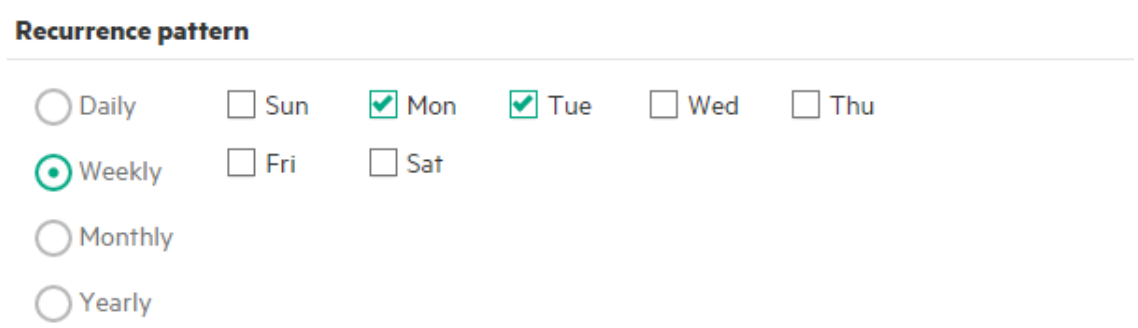	
Weekly	To run the flow on a weekly basis, select Weekly and then select the days of the week on which you want the flow to run. **Note:** If you select **Weekly**, you must select at least one day.
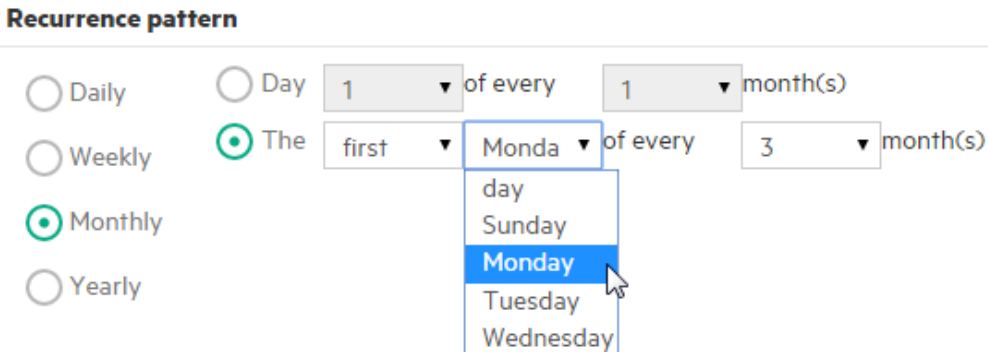	
Monthly – Day <n> of every month	To run the flow on a particular day of each month , select **Monthly**, select **Day** and then select the date. For example, select **15** to run the flow on the 15th of every month.

Monthly – The <nth> <day of week> of every <n> month(s)	To run the flow once a month on a particular day of the week, select the day of the week and whether it is the first, second, third, fourth, or last time that this day occurs in the month. For example, to run the schedule on the third Friday of each month, select **third**, then select **Friday**, and then select **1** month(s). To run the flow at intervals of several months, select a number of months. For example,to run a schedule once a quarter, on the first Monday of the quarter, select **first**, then select **Monday**, and then select **3** month(s). To run the flow on the last day of the month, select **last**, and then select **day** instead of a specific day of the week.
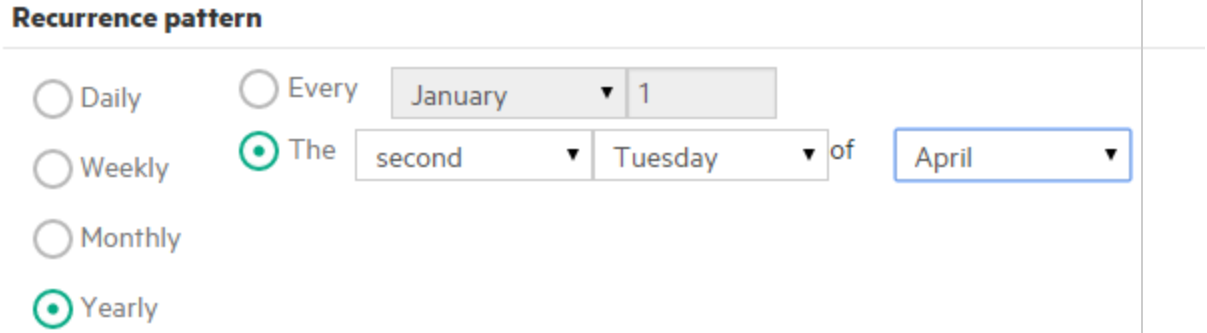	
Yearly – Every <month> <n>	To run the flow on a particular date each year, select **Yearly**, select **Every**, and then select the month and the date. For example, every July 4.
Yearly – The <nth> <day of week> of <month>	To run the flow once a year on a particular month and on a particular day of the week, select **Yearly**, select the month, select the day of the week and whether it is the first, second, third, fourth, or last time that this day occurs in the month. For example, the second Friday of each November.
Range of recurrence section	In the **Range of recurrence** section, you define when the schedule starts and ends.

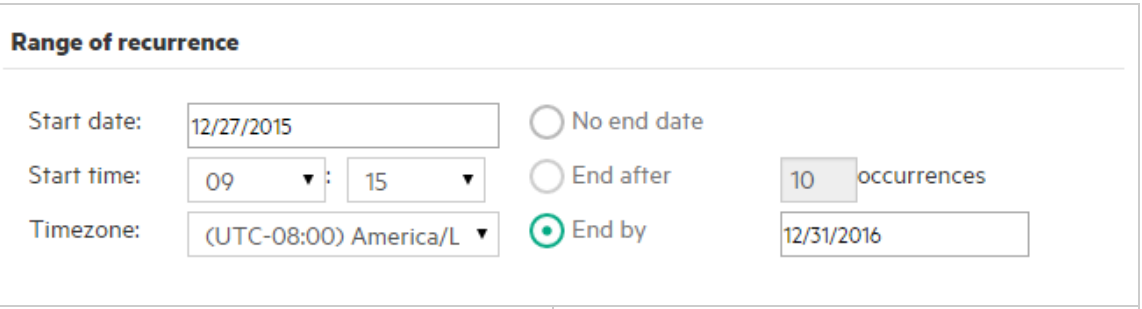

Start date	Enter the date on which you want the schedule to start. **Note:** If a schedule with an earlier start date undergoes content upgrade, the start date is changed to the current date.
Start time	Enter the time of day at which you want the schedule to start.
No end date	Select if you want the schedule to run continuously.
End after <n> occurrences	To limit the schedule to a specific number of runs, select **End after** and enter the number of runs after which the schedule will stop. This option is only available for daily schedules that have been set to run at intervals of a number of hours or minutes.
End by <date>	To end the schedule on a specific date, select **End by** and enter the date after which schedule will stop.
Timezone	From the **Timezone** list, select the time zone where you are located. **Note:** Time zones around the world are expressed as positive or negative offsets from Coordinated Universal Time (UTC). When you create a schedule, a time zone is selected by default from the alphabetically-ordered list that is currently available. HPE OO selects the first time zone from the list that has the offset equal to the client offset (computed with daylight savings, if relevant). Note that this timezone might not be geographically correct. For example, if you are located in a location where the time zone is UTC-09.00, the default time zone will be

	America/Anchorage. If desired, you can select the geographically correct location from the **Timezone** list. 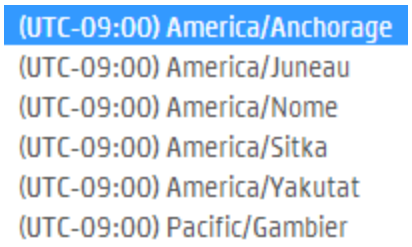
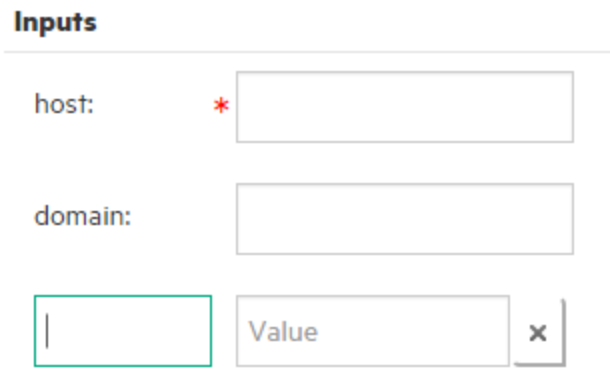	
Inputs	If the flow requires inputs, enter them in the **Inputs** section. This section will look different, depending on whether inputs are required.
Save	Click to save the schedule.

Tracking and Managing Flow Runs

 While a flow is running, you can track the run and monitor its progress in the **Run Explorer** tab. The table displays information for basic monitoring of your runs and those of other users.

The Run Explorer is auto-refreshed. The view is updated whenever a new run is started or if the run status, duration, or user is updated.

> **Note:** The name that appears under **Run Name** in the Run Explorer is not always the flow name:
>
> - If you gave the run a name, this is the name that appears.
> - For a scheduled run, if you gave the schedule a name, this is the name that appears.

You can use the Run Explorer to manage flow runs when there are multiple flows running. You can perform actions on one or more flow runs, such as pausing, resuming, reassigning, and canceling them.

The display of runs depends on the content permissions that your user has been assigned. You can only see runs of flows you have permission to view.

The actions that you can perform on flow runs depends on permissions that you have been assigned in your role:

- If you have been assigned a role with the **Manage Others' Runs** permission, you will be able to perform actions on the runs triggered by all users.
- If you have been assigned a role without this permission, you will be able to perform actions on the runs that you have triggered.

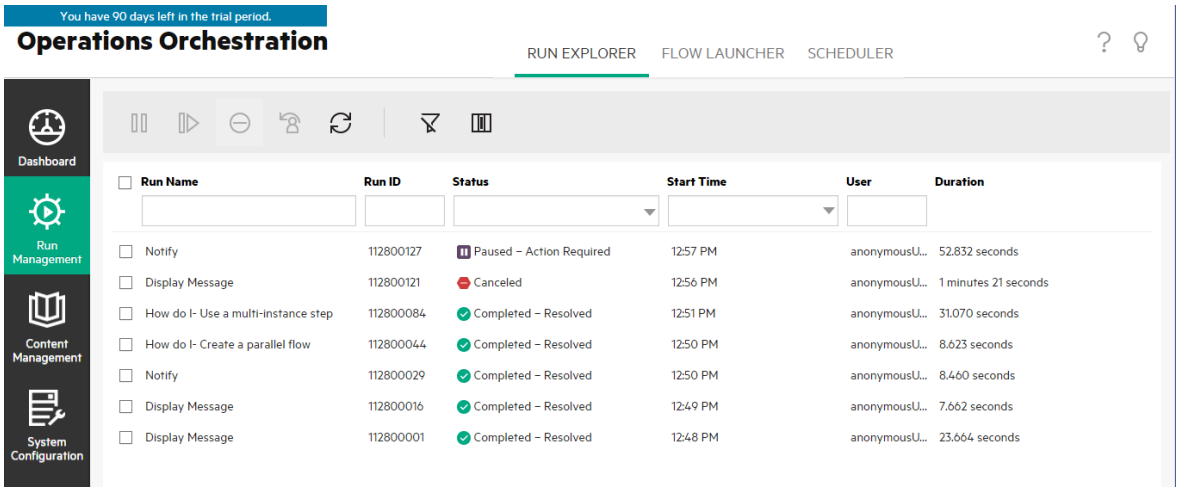

If you are unable to locate a flow run, you can use the filter capabilities of the Run Explorer to find the flow run that you need. For more information about filtering, see "Adjusting the Display of Flow Runs" on page 223.

What do you want to do?

View basic details about a flow run

1. Click the **Run Management** button to display the Run Management workspace.
2. Click the **Run Explorer** tab.
3. View the basic details about each run, in the **Runs** table.

Select multiple runs

You can perform the following actions on multiple runs: Cancel, Pause, Reassign, Resume.

There are two ways to select multiple runs in the Run Explorer:

- Click in the check box next to each run that you want to select
- Hold down the **Shift**/**Ctrl** keys and click anywhere in the row to select a run:
  - Use **Shift** to select multiple runs that are adjacent. Select the first run, and hold the **Shift** key while you select the last run.
  - Use **Ctrl** to select multiple runs that are not adjacent.

Display advanced details about a flow run

1. Click the **Run Explorer** tab.
2. Select the flow run that you want to look at in detail.
3. Click the **Drill down** button at the end of each row to display more information about the selected run.

   > **Note:** Alternatively, you can double-click a row to drill down to see this information.

   For more information, see "Testing and Troubleshooting a Flow Run" on page 229.

   > **Note:** The **Drill down** button is only displayed when you select a row in the table.

Pause a flow run

1. Click the **Run Management** button to display the Run Management workspace.
2. Click the **Run Explorer** tab.
3. Select one or more flows that are currently running.
4. Click the **Pause** button in the Run Explorer toolbar.

   > **Note:** The **Pause** button is only available if all selected runs are currently running.
   >
   > Note also that a run with **Pending pause** status will be timed out if the timeout threshold is reached.

The run status is changed to **Pending pause** and then to **Paused**.

Resume a flow run that was paused

You can resume a run that has **Paused – User Paused** status. If you have been assigned a role with the **Manage Others' Runs** permission, you can resume runs that were paused by others. Otherwise, you can resume runs that you paused.

1. Click the **Run Management**> **Run Explorer** tab.
2. Select one or more runs that were paused.
3. If the **Resume**  button is available in the Run Explorer toolbar, click it to resume the run (or runs).

If the **Resume** button is not available in the Run Explorer toolbar, this means that you need to interact with the system in order to complete these runs individually:

1. Click the **Drill down**  button at the end of the row to display more information about a run.
2. Click the **Expand**  button to expand the prompt message.
3. Enter the required information.
4. Click **Resume**  .
5. Repeat for the next run that needs to be resumed.

> **Note:** The **Resume** button is only available if all selected runs are currently paused.

Cancel a flow run

1. Click the **Run Management**> **Run Explorer** tab.
2. Select one or more flows that are currently running.
3. Click the **Cancel**  button in the Run Explorer toolbar.
4. In the confirmation dialog box, click **Yes** to confirm that you want to cancel the flow run (or runs).

> **Note:** The **Cancel**  button is only available if all selected runs can be canceled: if they are running or paused but not finished.

Reassign ownership of a flow run

Each run has an owner. By default, this is the person who started the run, but you can change the

owner during execution by assigning the run to a different user.

> **Note:** The owner of a run is the user who appears in the **User** column.

You can reassign multiple runs at the same time.

Reassignment is usually triggered as a result of a gated transition, where the run is paused because it doesn't have permission to proceed, or as a result of a handoff, where the author marked a transition as needing to be handed off to another user. However, you can also reassign a flow at any point in the flow run. If the flow is running, it will become paused so that it can be reassigned.

> **Note:** Reassignment does not validate that the reassigned user exists.

> **Note:** If you reassign a flow with the status of **Paused – Action Required**, note that the assigned user will not be able to resume it via the Run Explorer. The assigned user will need to drill down into the run and resume it from there.

1. Click the **Run Management**> **Run Explorer** tab.
2. Select one or more flows that are currently running or paused.
3. Click the **Reassign** button.
4. In the Reassign a Run dialog box, enter the user name of a different user and the domain to which this user belongs.
5. Click **Reassign**.

> **Note:** The **Reassign** button is only available if the selected runs are running or paused.

Hand off a flow run

Some flows contain gated transitions, requiring the current user to hand off the run to another user. The run pauses, and a dialog box enables the current user to send an email to another user, asking them to continue the run.

1. Click the **Run Management** > **Run Explorer** tab.
2. Select a run that is currently paused at the gated transition.
3. Click the **Drill down** button at the end of the row to display more information about the selected run.
4. Click the **Expand** button to expand the prompt message.

5. In the prompt message, click the **Hand off to another user (email this run)** link to create an email message to send to the user who needs to complete the flow run.
6. An email message opens, with a message for the person to whom you are handing off the flow run and a link to the flow run. Enter the email address of this person and send the message.

Refresh the runs in the Run Explorer

The **Refresh** button refreshes the view that is currently below in the Run Explorer—whether this is the run list or the drill down. For information about the drill down, see "Testing and Troubleshooting a Flow Run" on page 229.

- Click the **Refresh** button in the Run Explorer toolbar.

> **Note:** It is also possible to refresh the page by pressing F1 on your keyboard, but this loads the entire page. Clicking the **Refresh** button only refreshes the data display, so this is much faster.

Reference Material

Run Explorer

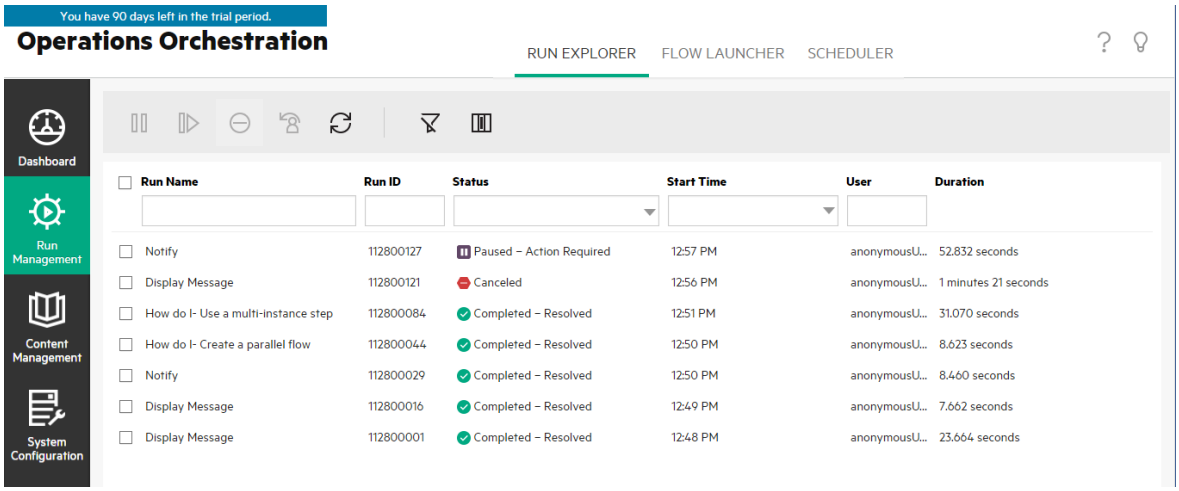

GUI item	Description
Selection check box	Click the check box next to one or more runs, in order to select them.

Pause  **button**	Click to pause the selected run or runs. The run will stop until it is resumed. This button is only available for runs that are currently running.
Resume  **button**	Click to resume the selected run (or runs). This button is only available for runs that are paused. **Note:** If a flow with the status of **Paused – Action Required** is reassigned, the assigned user will not be able to resume it via the **Resume**  button in the Run Explorer. The assigned user will need to drill down into the run and resume it from there.
Cancel  **button**	Click to cancel the selected run or runs.
Refresh  **button**	Click to refresh the runs that are displayed in the Run Explorer.
Reassign  **button**	Click to reassign the selected run or runs to a different user. **Note:** If you reassign a run with the status of **Paused – Action Required**, the assigned user will not be able to resume it via the **Resume**  button in the Run Explorer. The assigned user will need to drill down into the run and resume it from there.
Clear Filter  **button**	Click to remove the filters and display all the runs.
Filter boxes	In the filter boxes at the top of the columns, enter text to filter the displayed runs.
Run Name	Displays the name of the run. This is not always the flow name: • If you gave the run a name, this is the name that appears. • For a scheduled run, if you gave the schedule a name, this is the name that appears.
Run ID	Displays a unique Run ID for the run.
Status	Displays the current status of the flow run. Options include:  **Completed – Resolved**  **Completed – Error**  **Completed – No Action Taken**

	 **Completed – Diagnosed**  **Completed – Custom**: for CloudSlang flows only. For more information, see "Working with CloudSlang Content Flows" on page 120.  **Failed to Complete**: because of a Run Management exception  **Canceled – Stopped by User**  **Canceled – Timeout Expired**: the run has exceeded the maximum duration  **Running**  **Paused – User Paused**  **Paused – Action Required**: the run will not continue until the action is completed. For example, enter missing information, or change to a user with permission to run the next step.  **Paused – No Workers in Group**: there is a worker group that has not yet been mapped to any workers. Before you can resume the run, map the worker group to actual workers. See "Setting Up Configuration Items for a Content Pack" on page 164.  **Pending Pause**: the user has asked to pause the run, but the run has not been paused yet.  **Pending Cancel**: the user has asked to cancel the run, but the run has not been canceled yet.
User	Displays the current owner of this run. This may be the user who triggered the run, the user who last modified the run, or the user to whom the run was reassigned. If this last user is logged in as an anonymous user, the **User** column displays **anonymousUser**.
Start Time	Displays the time at which the run started.
Duration	Displays the duration that the run has been running. To refresh this value, press F5.

Drill Down button	Click to drill down and display more information about the flow run.

Adjusting the Display of Flow Runs

Adjusting Filters

The Run Explorer can display up to 200 flow runs in a single page. So if you have many flows running at the same time, it is recommended to use filters to locate the flow that you need.

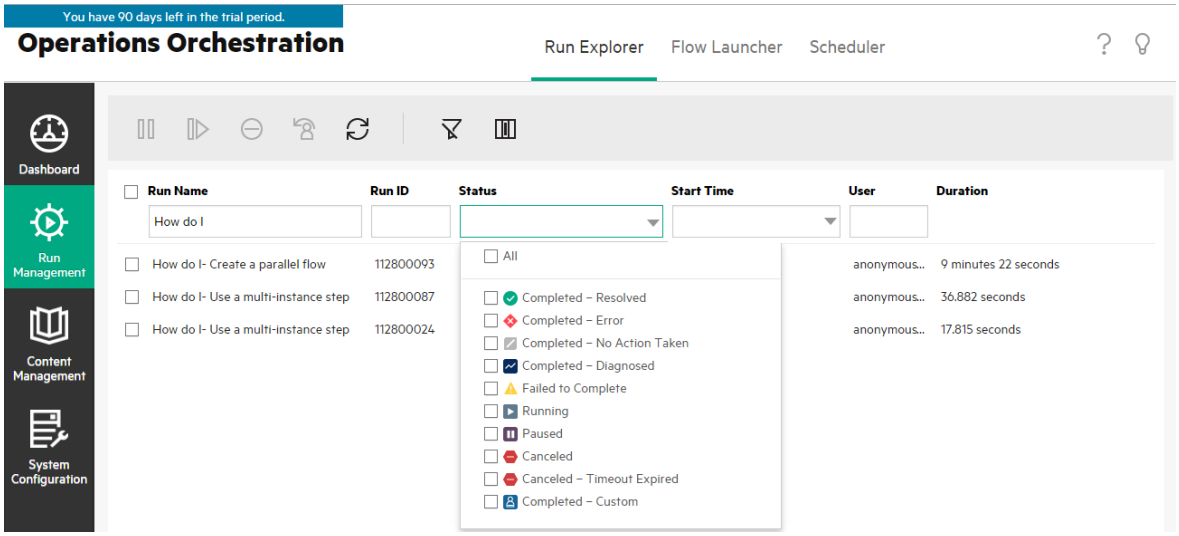

You can create a filter to view runs that contain a specific word in their name, runs with a specific status, runs that were started by a specific user, or runs that started within a defined time range.

You can apply multiple filters at the same time. For example, you can filter the Run Explorer to display only failed runs with the word "Ping" in their name, which were started on the first of June.

To apply a filter, select or enter text in the filter boxes in the Run Explorer toolbar. Once you have entered the criteria, the filter is applied—there is no need to click a button to apply it.

It is also possible to filter runs via API. For more information, see the *HPE OO API Guide*.

Adjusting Columns

You can also adjust the width and position of a column, and decide which columns to show and which to hide.

Changes that you make to these settings (column width, column order, hidden columns, and applied filter) are saved in the local cache of the browser.

> **Note:** Changes that you make to the columns will affect all users on this machine, regardless of which user is logged in.

> **Note:** When you hide a column, if you were using a filter, this deletes the filter.

What do you want to do?

Create a filter according to run status

1. Click the **Run Management** button to display the Run Management workspace.
2. Click the **Run Explorer** tab.
3. From the **Status** list, select the status that you want to filter by. You can select more than one status.

   The Run Explorer only displays runs with one of the selected statuses.

   > **Note:** After you have selected a status for the filter, the name of this status appears as the title of the **Status** list.

   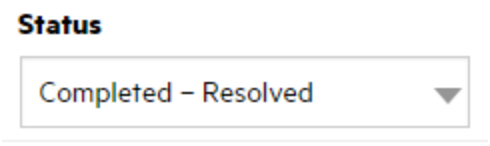

Create a filter according to start time

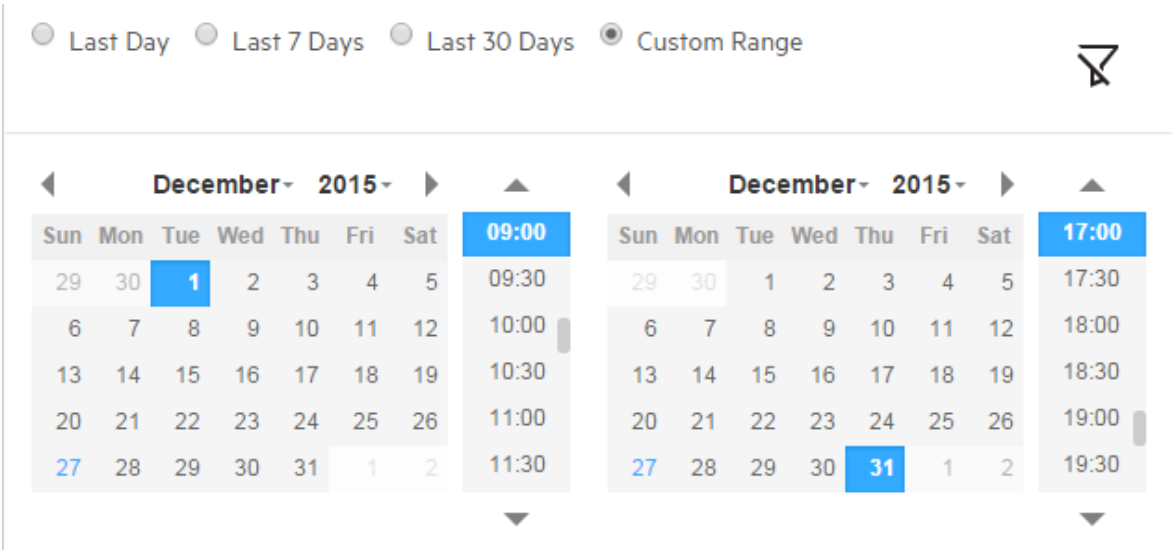

1. In **Run Management**, click the **Run Explorer** tab.
2. In the Run Explorer toolbar, select the arrow next to **Start Time**.
3. Select one of the following options:
   - **Last Day** – to display only runs that started within the last 24 hours.
   - **Last 7 Days** – to display only runs that started within the last week.
   - **Last 30 Days** – to display only runs that started within the last 30 days.
   - **Custom Range**, and then select the start and finish of the time range that you want to filter by. You can select both the date and the time of day. The calendar on the left is for the start of the time range and the calendar on the right is for the end of the time range.

     For example, you may want to display only runs that started between 9:00 and 10:00 AM on Monday August 20, 2013.

Create a filter according to user name

Filters are not case-sensitive, so you can type in upper or lower case.

1. In **Run Management**, click the **Run Explorer** tab.
2. In the **User** text box, enter part or all of the user name.

   The Run Explorer displays only runs that were started by the specified user. For example, entering your user name displays only the flow runs that you have run.

Create a filter according to run name

Filters are not case-sensitive, so you can type in upper or lower case.

1. In **Run Management**, click the **Run Explorer** tab.
2. In the **Run Name** text box, enter part or all of the run name.

   The Run Explorer displays only runs that contain this text in their name. For example, typing “network” returns all the rows where the filtered field contains the word “network”.

Create a filter according to run ID

1. In **Run Management**, click the **Run Explorer** tab.
2. In the **Run ID** text box, enter part or all of the run ID.

   The Run Explorer displays only the run with this run ID.

Remove a filter

1. In **Run Management**, click the **Run Explorer** tab.
2. To remove an individual filter, clear the relevant filter text box.
3. To remove all the filters, click the **Clear Filter** >  button.

   The filters are no longer applied.

Change the position of a column

1. In **Run Management**, click the **Run Explorer** tab.
2. In the Run Explorer toolbar, drag a column name to the left or the right, to adjust the column's position in the Run Explorer.

Adjust the width of a column

1. In **Run Management**, click the **Run Explorer** tab.
2. In the Run Explorer toolbar, drag the edge of a column name, to adjust its width.

Hide/display a column

1. In **Run Management**, click the **Run Explorer** tab.
2. In the Run Explorer toolbar, click the **Select Columns** button to display the column picker.

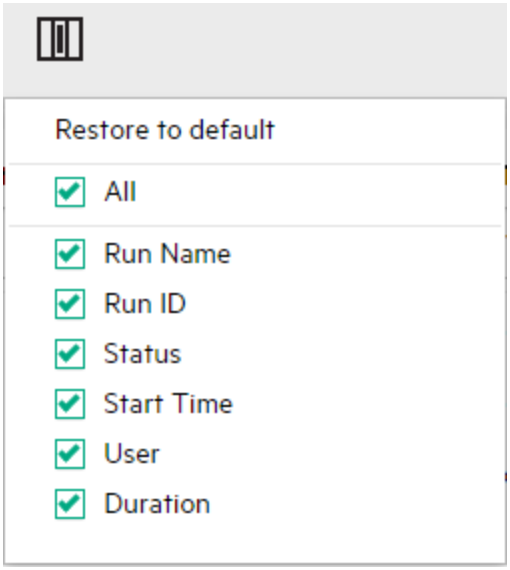

3. To hide a column, clear the check box next to its name.

Reference Material

Run Explorer toolbar - Filter options

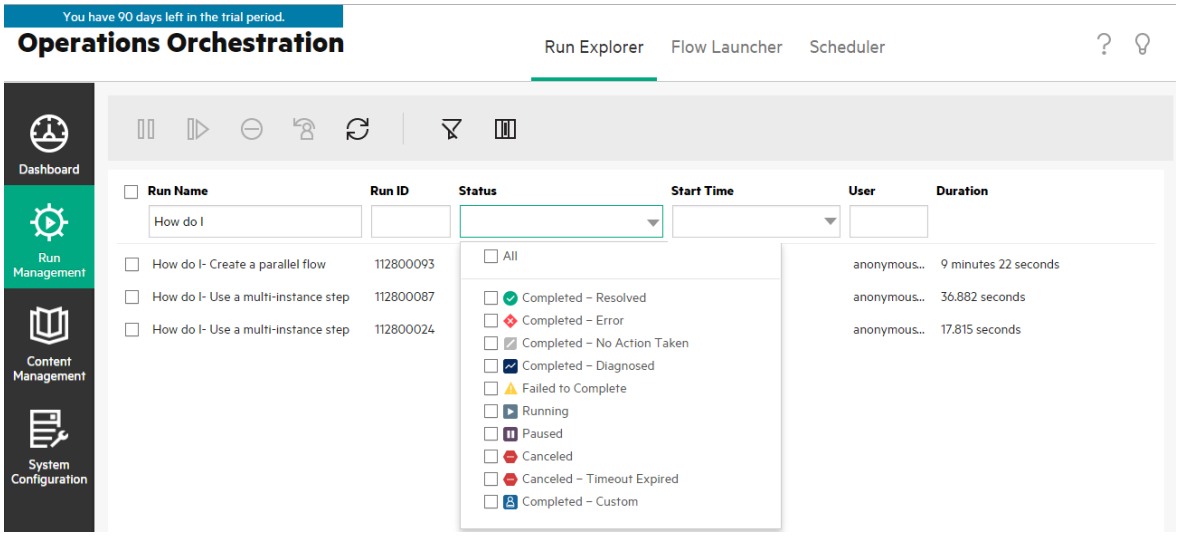

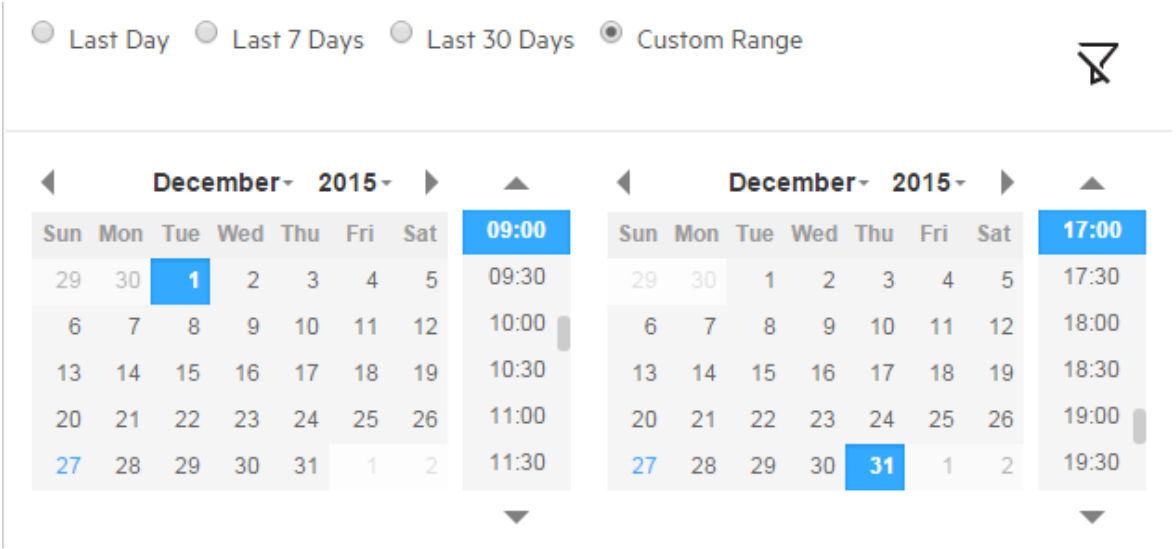

GUI item	Description
Filter by: Status	From the **Status** list, select the status or statuses that you want to filter by.
Filter by: User	Enter part or all of the user name.
Filter by: Start Time	Select the arrow next to **Start Time** to filter according to the date and time that the runs were started. Select one of the options: • **Last Day** – to display only runs that started within the last 24 hours. • **Last 7 Days** – to display only runs that started within the last week. • **Last 30 Days** – to display only runs that started within the last 30 days. • **Custom Range**, and then select the start and finish of the time range that you want to filter by. You can select both the date and the time of day. The calendar on the left is for the start of the time range and the calendar on the right is for the end of the time range. For example, you may want to display only runs that started between 9:00 and 10:00 AM on Monday August 20, 2013.
Filter by: Run Name	Enter part or all of the run name.
Filter by: Run ID	Enter part or all of the Run ID.
Clear Filter button	Click to remove the filters and display all the flows.

Testing and Troubleshooting a Flow Run

After you have triggered a flow run, you can drill down into the run and investigate its progress in detail. This is particularly helpful if a flow run fails.

When you click the **Drill down** button at the end of a row in the Run Explorer, this displays detailed information about the selected run.

> **Note:** Alternatively, you can double-click a row to drill down to see this information.

To the left, the Run Tree displays the steps in the run and the transition messages. The Run Tree highlights the currently running step and displays information about the progress of the run. If a flow fails, you will see at which step the failure occurred. The steps and transition messages are automatically refreshed as the run progresses. While a flow is running, if you click a step, the auto-refresh will pause. In order to resume you need to click the **Resume Refresh** button.

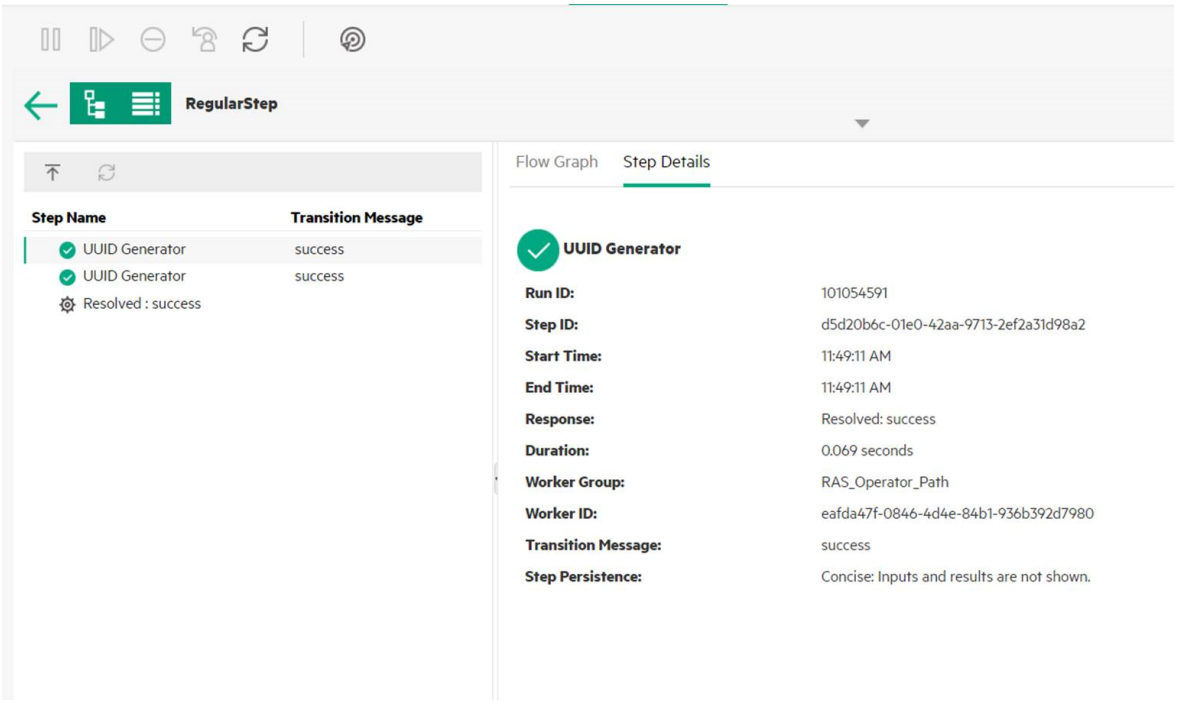

There are a number of collapsible views that you can use to display different kinds of information:

- **Information about the flow**—UUID, flow name, path, description, flow inputs, flow outputs, and so on.

  Click the **Down** arrow  in the center of the Run Explorer drill down view toolbar.

- **Information about a selected step in the flow**—start time, end time, step result, duration, inputs, results, worker group, the ID of the worker that executed the step, and so on. Click the **Step Details** tab.

  

  The **Step Details** tab displays details about the step that is currently running. If you select a step in the Run Tree, the **Step Details** tab displays details about the selected step.

  The Step Details tab may contain Concise or Detailed information about the step. The information shown depends on the option selected for the step in the Advanced tab of the Step Inspector. See "**Step Inspector > Advanced tab**" in the *HPE OOStudio Authoring Guide* for details.

  > **Note:** If any of these details were defined as sensitive data, they will be encrypted and will appear as asterisks.
  >
  > All sensitive data is encrypted in-memory and in the database during the run and later. The event is persistent and is presented as asterisks in the UI.
  >
  > This property is propagated between elements. For example, if a step result was marked as sensitive data, the step result passes this behavior to every input/step result that is assigned from it.
  >
  > This includes sensitive data returned by a scriptlet.
  >
  > For information about how to define data as sensitive, see "Creating Outputs and Results" in the *HPE OO Studio Authoring Guide*.

  > **Note:** You can drag the slider on the left edge of the **Flow Info** pane to adjust its width.

- **A graphical view of the flow**, with the currently running step highlighted. Click the **Flow Graph** tab.

  The Flow Graph highlights the step that is currently running. If you select a step in the Run Tree, the Flow Graph highlights the selected step.

  > **Note:** You can adjust the way that the flow graph is displayed, by dragging it across the screen to view hidden item, or by using the zoom buttons to adjust the size of the image.

- **The Run Log**—a table with information about the entire run (steps, start time, duration, inputs, results, and so on), with the option to find a specific step.

  Click the **Run Log**  button to display the **Run Log** pane.

  > **Note:** If the run includes subflows, the steps in the subflows are indented in the run log.

  Click the **Run Tree**  button to collapse the **Run Log** pane and show the Run Tree.

  > **Note:** You can click the  button in the toolbar to toggle between viewing a selected step in the Run Tree and in the Run Log.

> **Note:** There are also cases where a run has no drill down details:
>
> - If a run comes from a content upgrade from version 9.x.
> - If the run details were cleared in the system.

What do you want to do?

Display details about the steps in a flow run

1. Click the **Run Management** button to display the Run Management workspace.
2. Click the **Run Explorer** tab.
3. Select the flow run that you want to look at in detail.
4. Click the **Drill down**  button at the end of the row to display the drill down information about the selected run.

   To the left, the Run Tree displays the steps in the run and the transition messages. The Run Tree highlights the currently running step and displays information about the progress of the run. If a flow fails, you will see at which step the failure occurred.

   > **Note:** The steps and transition messages are automatically refreshed as the run progresses. While a flow is running, if you click a step, the auto-refresh will pause. In order to resume you need to click the **Resume Refresh**

 button.

5. When you click a step in the Run Tree, you display information about that step in the **Run Info** pane on the right.

   In the **Run Info** pane, click the **Step Details** tab to display information about the currently running step, such as the start time, duration, inputs, results, description, worker group, the ID of the worker that executed the step, and so on.

   If you select a step in the Run Tree, the **Step Details** tab displays details about the selected step.

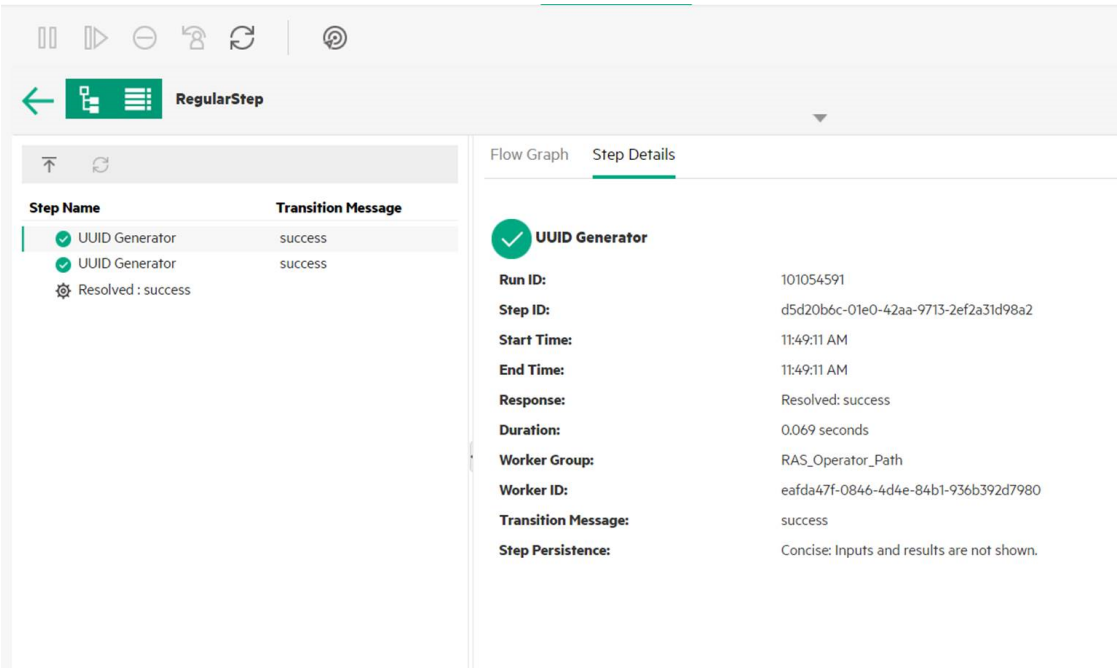

> **Note:** The fields displayed in this section will vary between different steps, depending on whether they include inputs, results, and so on.

6. (Optional) You can click the  button in the toolbar to view the selected step in the Run Log.
7. Click the **Back** button to return to the main **Run Explorer** window.

Display the graph of a flow run

You can display a graphical image of a flow run, and highlight the position of the currently running (or selected) step.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at in detail, and click the **Drill down** button at the end of the row.
3. Click a step in the Run Tree on the left.
4. In the **Run Info** pane on the right, click the **Flow Graph** pane to display a graphical view of the flow, with the currently running step highlighted.

   If you select a step in the Run Tree, the Flow Graph highlights the selected step.

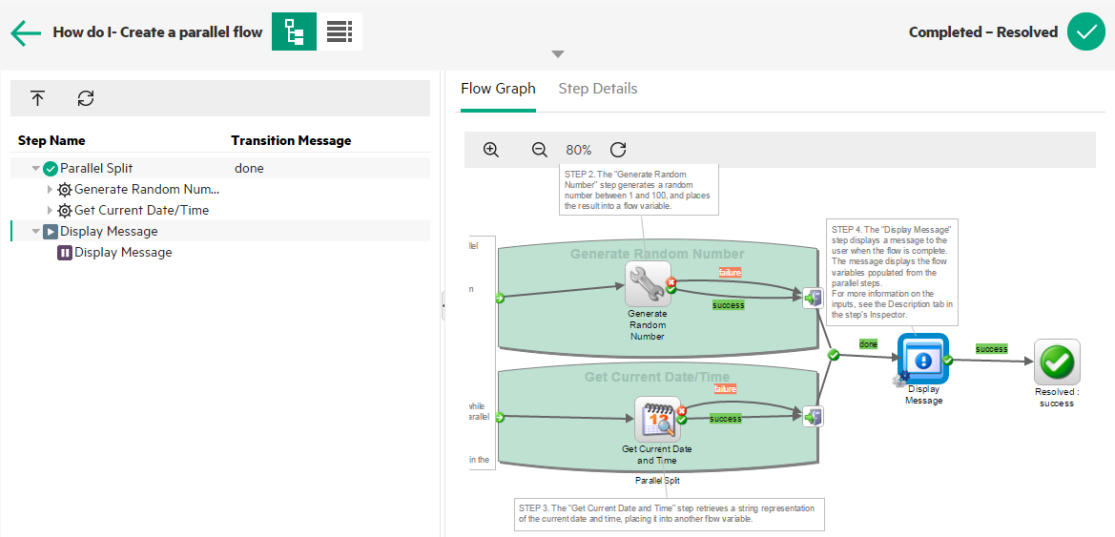

5. You can adjust the way that the flow graph is displayed:
   - Drag the flow graph across the screen, to display elements that are not visible.
   - Click the zoom buttons to increase or decrease the size of the flow graph.
   - Click the **Reset** button to reset the view to the default size and position.

Display the Run Log with details of the entire run

The Run Log is a table with information about the entire run (steps, start time, duration, inputs, results, and so on). This information is useful for identifying issues about a run, while troubleshooting.

The Run Log displays a maximum of 500 runs in one page, in order to meet performance requirements. Runs with over 500 steps are displayed in multiple pages. You can use the paging functionality to navigate through the pages.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Click the **Run Log** button to display the **Run Log** pane, with details about the entire run.

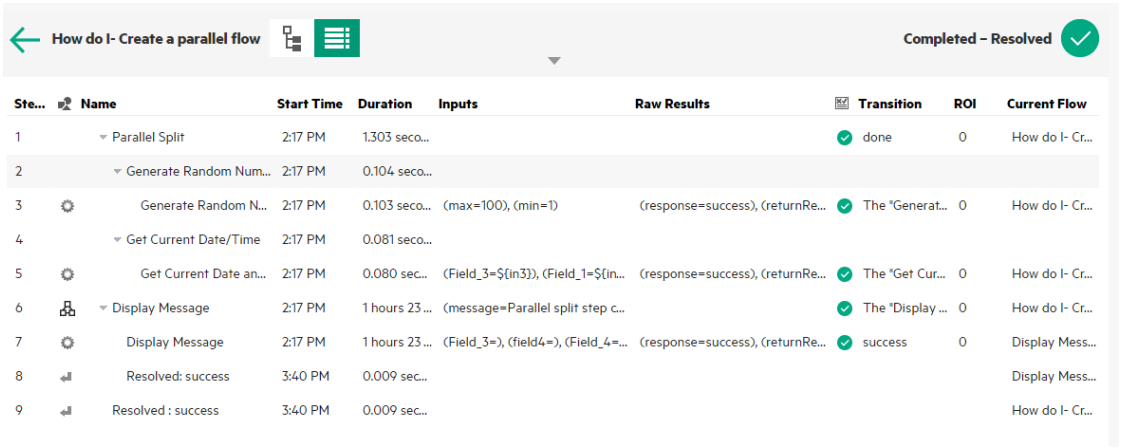

> **Note:** If the run includes subflows, the relevant steps are indented.

4. For runs of over 500 steps, which are displayed in multiple pages, navigate through the pages:
   - Click the arrow buttons
   - Enter the number of a specific page

5. To hide or display columns, click the **Select Columns** button to display the column picker. Select the check box next to a column name to display that column or clear the check box to hide it.

> **Note:** When you hide a column, this does not affect the results if you run a search.

6. You can click the **Run Tree** button to collapse the **Run Log** pane and show the Run Tree. Or click the button in the toolbar to view the selected step in the Run Tree.

Find a step in the Run Log

In a long run, it may be difficult to find a particular step. Use the **Find** button to help locate the step that you want to see.

> **Note:** Searching in the Run Log is only supported if the storage space taken up by the searched field is no more than 4,000 bytes. Note that Japanese and Chinese take up about 3 bytes per character, while other (non-English) languages take up about 1.5 bytes per character. Also note that in the case of transition descriptions, the 4,000 bytes are shared among all of the locales in which the description is available.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Click the **Run Log** button to display the **Run Log** pane.
4. Click the **Find** button to display the Find dialog box.

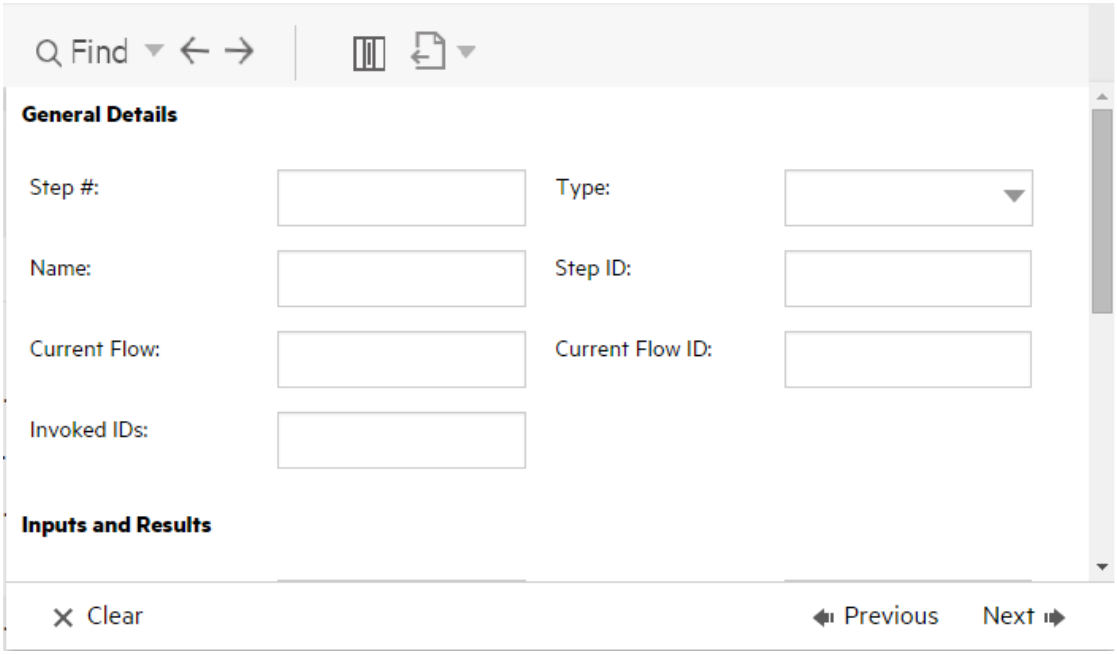

5. Enter the search criteria, scrolling down to see those that are hidden.

> **Note:** This is not a filter. Steps that correspond to the search criteria are highlighted, but the unselected steps are displayed, so that you can see the context of the highlighted steps.

6. Click **Next** to highlight, in turn, each step that corresponds to the search criteria.
7. If desired, you can close the Find dialog box and use the **Next** and **Previous**  arrows in the toolbar to navigate across the search results.

Toggle between viewing a step in the Run Tree + Flow Graph and in the Run Log

You can move in either direction:

- While a step is selected in the Run Tree, you can jump to that same step in the Run Log to display more information about that step.
- While a step is selected in the Run Tree, you can jump to that same step in the Run Log and Flow Graph, to see how the step fits into the overall picture of the flow.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down**  button at the end of the row.
3. When you select a step in the Run Tree, click the  button to display the Run Log with the selected step highlighted.
4. When you select a step in the Run Log, click the  button to display the Run Tree with the selected step highlighted. The selected step is also highlighted in the Flow Graph.

Collapse the Run Tree

If the Run Tree is expanded, you can collapse the expanded steps.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down**  button at the end of the row.
3. Click the steps in the Run Tree, in order to expand them, and display subflows and lanes.
4. Click the **Collapse All**  button to collapse the Run Tree and display only the top level of the steps.

Adjust the width of the Run Info pane

You can use the slider at the edge of the **Run Info** pane to adjust its size.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Drag the edge of the slider to adjust the width of the **Run Info** pane.
4. Use the open/close button in the slider to toggle between expanding and collapsing the pane.

Display Details about the flow

While drilling down to a flow run, you can view details about the original flow, such as flow UUID, flow name, path, description, flow inputs, flow outputs, and so on.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Click the **Down** arrow to the right of the flow run name, to expand the **Flow Info** pane.

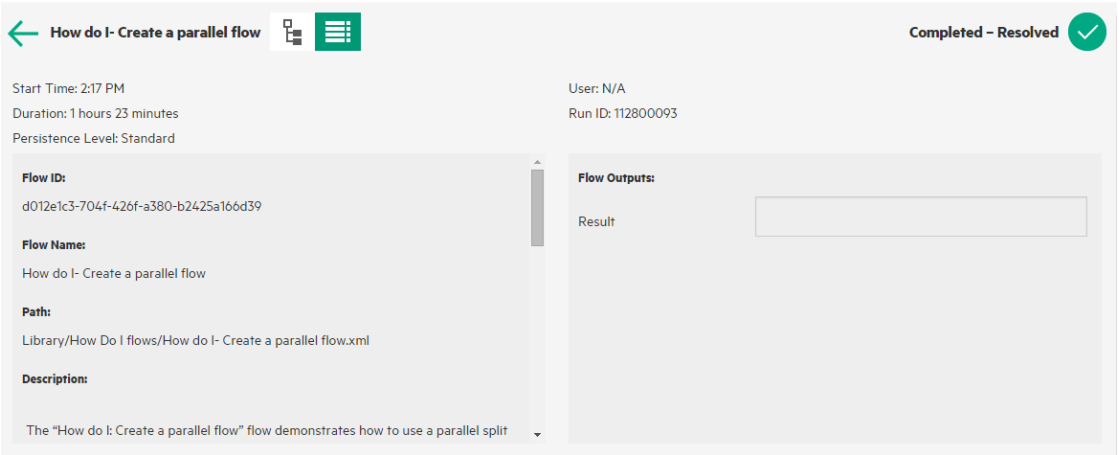

4. If you can't see all the text in the pane, use the slider to display the text that is lower down in the

pane.

5. Click the **Up** arrow to collapse the **Flow Info** pane.

Export the Run Log as a text file

You can export the Run Log into a text file, listing all the steps' logs, which occurred in the flow run. This can help identify the causes of a problem. It may be useful to send the this file to the flow author, when a flow needs fixing.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Click the **Run Log** button to display the **Run Log** pane.
4. Click the **Export** button and from the dropdown, select Export to Text .

   A button with the name of the text file appears at the lower edge of the HPE OO window. You can click it to open the file in your default text editor.
5. You can send the text file to the flow author, when requesting a change to a flow.

Export the Run Log as a CSV file

You can export the Run Log into a CSV file, listing all the steps' logs, which occurred in the flow run. This can help identify the causes of a problem. It may be useful to send the this file to the flow author, when a flow needs fixing.

> **Note:** In order to open the CSV file in Excel, make sure that Excel is configured correctly for character encoding. For more information, see http://superuser.com/questions/280603/how-to-set-character-encoding-when-opening-excel.

1. Click the **Run Management** button to display the Run Management workspace, and click the **Run Explorer** tab.
2. Select the flow run that you want to look at, and click the **Drill down** button at the end of the row.
3. Click the **Run Log**

 button to display the **Run Log** pane.

4. Click the **Export** button  and from the dropdown, select 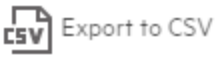 .

   A button with the name of the CSV file appears at the lower edge of the HPE OO window. You can click it to open the CSV file in Excel. 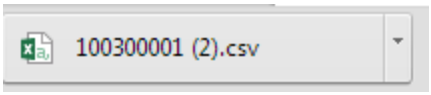

5. You can send the CSV file to the flow author, when requesting a change to a flow.

Reference Material

Run Explorer drill down view toolbar

When you drill down to see the details of a flow run, the toolbar at the top remains constant, and the details displayed below vary according to your selections.

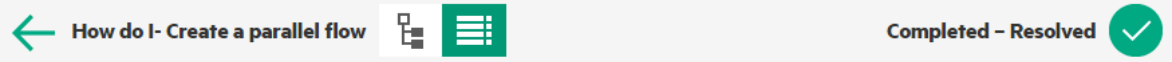

GUI item	Description
Back  **button**	Click to return to the main **Run Explorer** window.
Run Tree  **button**	Click to display the **Run Tree** pane, when the Run Log is displayed.
Run Log  **button**	Click to display the **Run Log** pane, when the Run Tree is displayed.
Run name	Displays the name of the flow run.
Down arrow 	Click to expand the **Flow Info** pane.
Status	Displays the status of the flow run:

	Success  or Error 

Run Tree

When you drill down to see the details of a flow run, the Run Tree appears on the left. The Run Tree displays all the steps in the flow run and their transition messages. The action toolbar at the top enables you to perform actions on the run.

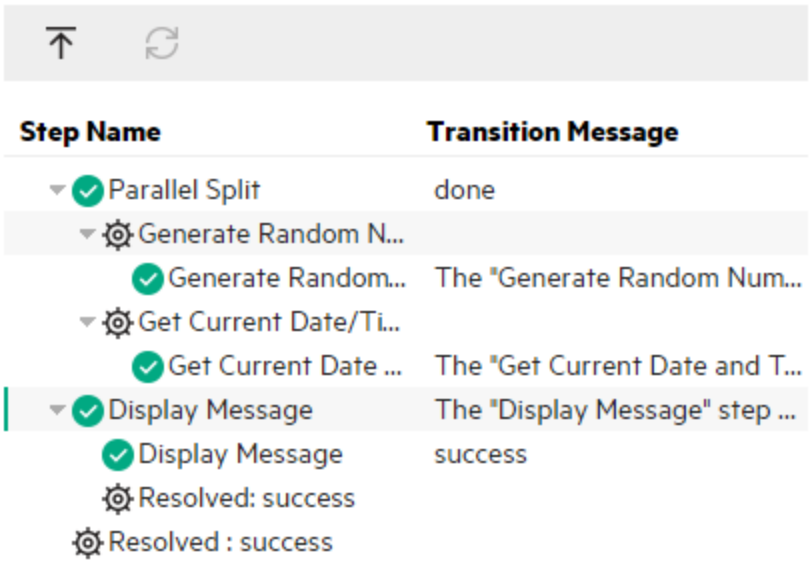

GUI item	Description
Action Toolbar	The action toolbar shows different buttons, depending on whether the Run Log or the Run Tree is displayed.
Pause  **button**	Click to pause the selected run. The run will stop until it is resumed. This button is only available for runs that are currently running.
Resume  **button**	Click to resume the run. This button is only available for runs that are paused.
Cancel  **button**	Click to cancel the run.
Reassign 	Click to reassign the run to a different user.

button	
Refresh  **button**	Click to refresh the run.
Go to Log Tree  **button**	While a step is selected in the Run Tree, jump to that same step in the Run Log.
Collapse All  **button**	Collapse the Run Tree and display only the top level of the steps.
Resume Refresh  **button**	While a flow is running, if you click a step, the autorefresh will pause. In order to resume you need to click the **Resume Refresh** button.
Step Name	Displays the name of each step in the flow run. Next to each step, an icon displays the step's status. Click a step to display its details in the **Flow Info** pane to the right.
Transition Message	Displays the transition message for the step, if one exists.

Run Info > Step Details tab (Concise Step Persistence)

Click the **Step Details** tab to display information about the currently running step. If you select a step in the Run Tree, the **Step Details** tab displays details about the selected step.

The fields displayed in the **Step Details** tab may vary between different types of step.

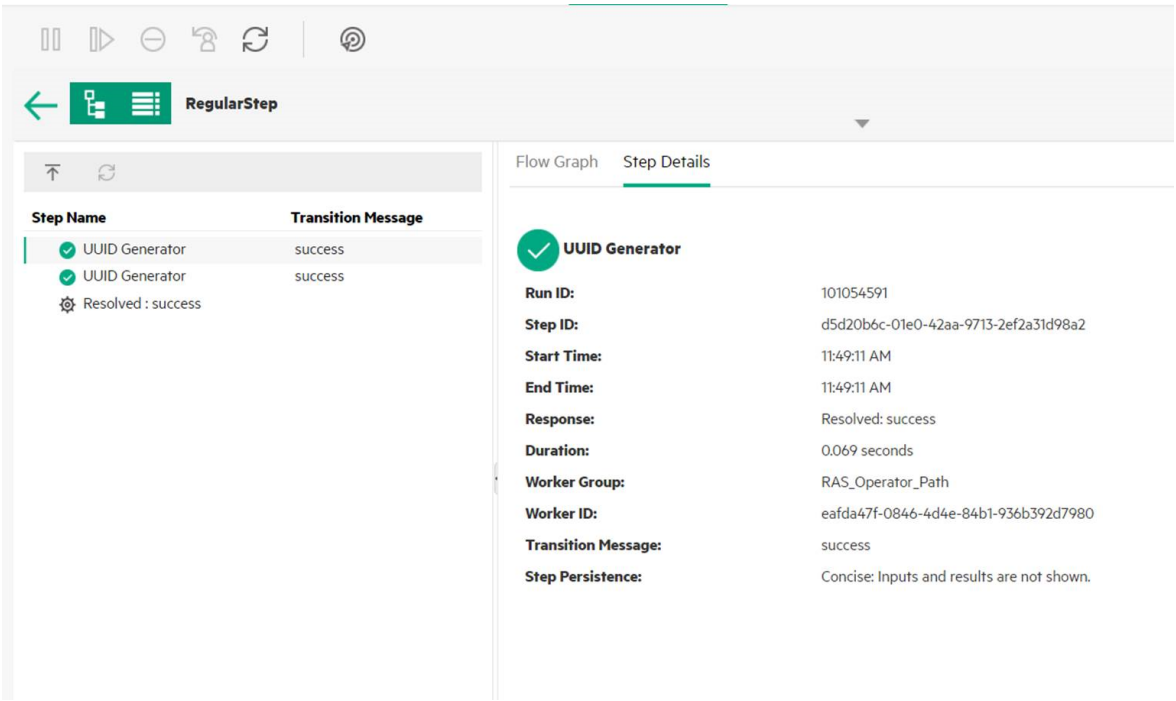

GUI item	Description
Run ID	Displays the UUID of the selected run.
Step ID	Displays the UUID of the selected step.
Start Time	Displays the time when the selected step started.
End Time	Displays the time when the selected step ended.
Response	Displays the response of the selected step, if one exists.
Duration	Displays the duration of the selected step.
Description	Displays the description of the selected step, if one exists.
Workers Group	Displays the worker group of the selected step, if one exists.
Worker ID	Displays the UUID of the worker that executed the selected step, if one exists.
Transition Message	Displays the message that appears during the transition of the selected step, if one exists.
Step Persistence	Displays the persistence level for this step as defined in the Advanced tab in the Step Inspector in Studio. For Concise, the following text appears: `Concise: Inputs and results are not shown`

Run Info > Step Details tab (Detailed Step Persistence)

Click the **Step Details** tab to display information about the currently running step. If you select a step in the Run Tree, the **Step Details** tab displays details about the selected step.

The fields displayed in the **Step Details** tab may vary between different types of step.

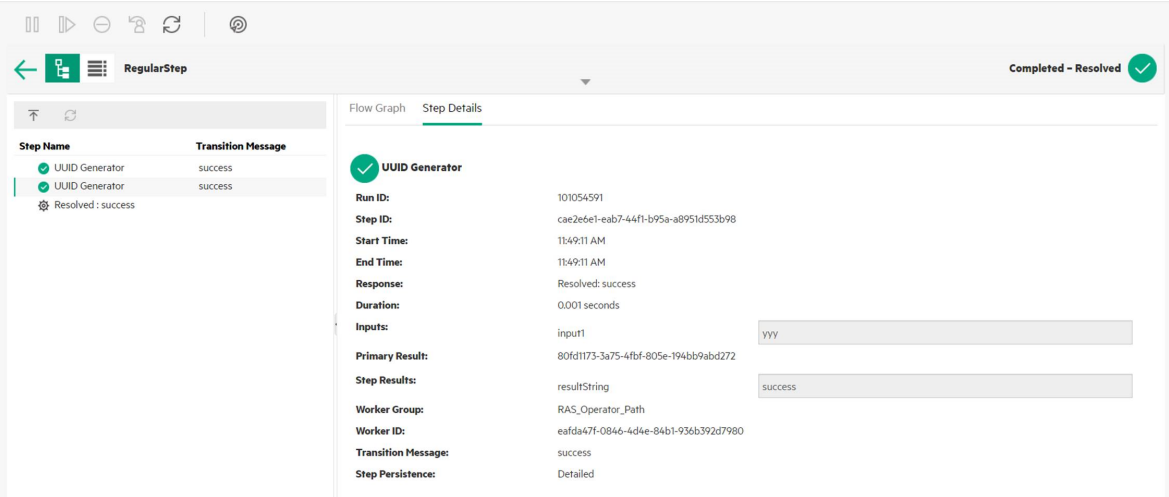

GUI item	Description
Run ID	Displays the UUID of the selected run.
Step ID	Displays the UUID of the selected step.
Start Time	Displays the time when the selected step started.
End Time	Displays the time when the selected step ended.
Step Response	Displays the response of the selected step, if one exists.
Duration	Displays the duration of the selected step.
Inputs	Displays the inputs of the selected step, if these exist. **Note:** If these inputs were defined as sensitive data, they will be encrypted and will appear as asterisks.
Primary Result	Displays the primary result of the selected step, if one exists. **Note:** If the raw result is marked as sensitive, then the primary result will be encrypted and will appear as asterisks.
Step Results	Displays the results of the selected step, if any exist. **Note:** If these results were defined as sensitive data, they will be encrypted and will appear as asterisks.

Worker Group	Displays the worker group of the selected step, if one exists.
Worker ID	Displays the UUID of the worker that executed the selected step, if one exists.
Transition Message	Displays the message that appears during the transition of the selected step, if one exists.
Step Persistence	Detailed.

Run Info > Flow Graph tab

Click the **Flow Graph** tab to display a graphical view of the flow, with the currently running step highlighted. If you select a step in the Run Tree, the Flow Graph highlights the selected step.

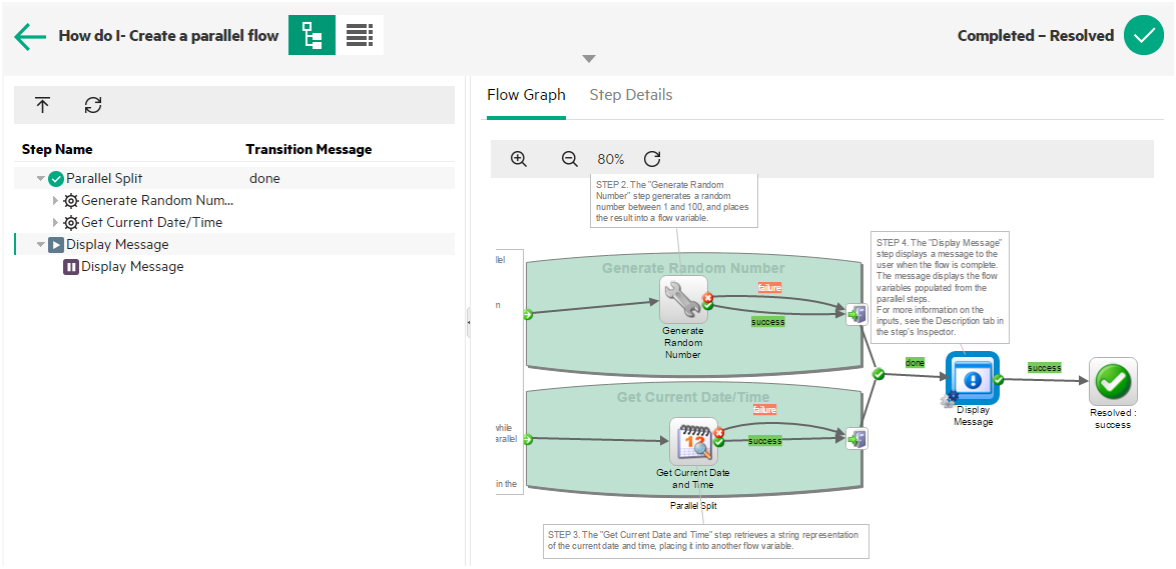

GUI item	Description
Zoom In button	Click to increase the size of the flow graph.
Zoom Out button	Click to decrease the size of the flow graph.
Reset button	Click to reset the graph to the default size and position.

Run Log

Click the **Run Log**

button to display the **Run Log** pane with details about the entire run. The action toolbar at the top enables you to perform actions on the run.

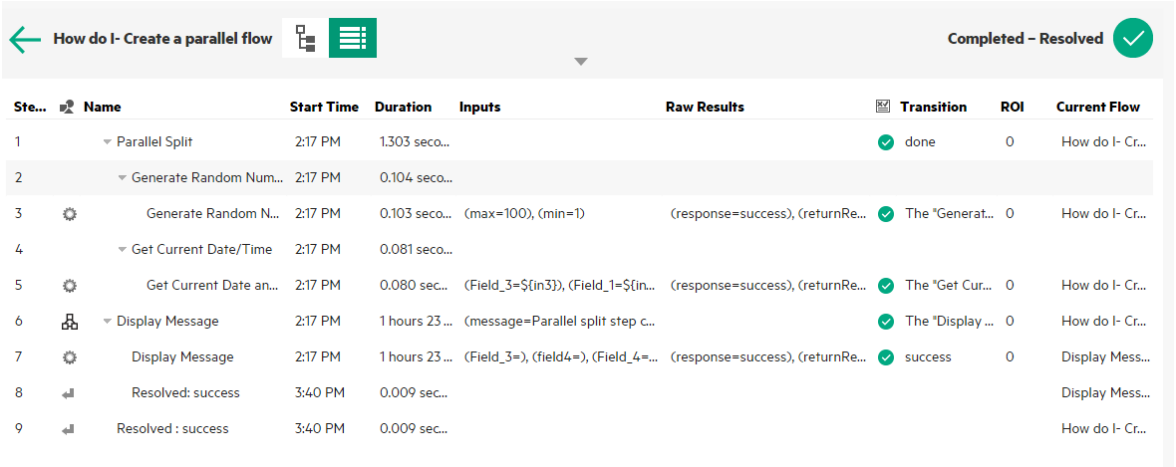

GUI item	Description
Action Toolbar	The action toolbar shows different buttons, depending on whether the Run Log or the Run Tree is displayed.
Pause  **button**	Click to pause the selected run. The run will stop until it is resumed. This button is only available for runs that are currently running.
Resume  **button**	Click to resume the run. This button is only available for runs that are paused.
Cancel  **button**	Click to cancel the run.
Reassign button 	Click to reassign the run to a different user.
Refresh  **button**	Click to refresh the run.
Go to Run Tree view  **button**	While a step is selected in the Run Log, jump to that same step in the Run Tree.
Find	Click to open the Find dialog box, to locate a step.
	After running a search, close the Find dialog box and use the arrows in the toolbar to navigate across the search results.

Export All to Text/CSV button 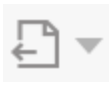	Click to export the run log into a text/CSV file, listing all the events that occurred in the flow run.
Step #	Displays the step number of each step
Icon 	Displays the type of step. Icons represent operation, subflow, return step, and other.
Name	Displays the name of the step. If the step is part of a subflow, the step name is indented.
Start Time	Displays the time when each step started.
Duration	Displays the duration of each step.
Inputs	Displays the inputs of each step, if these exist.
Raw Results	Displays the results of each step. **Note:** If these results were defined as sensitive data, they will be encrypted and will appear as asterisks.
Response	An icon represents the response of each step: Resolved, Error, Diagnosed, No Action Taken.
Transition	Displays the description of the outgoing transition, or its name if there is no description.
ROI	Displays the ROI of the outgoing transition.
Current Flow	Displays the name of the flow containing the step. If there are subflows, this contains the name of the subflow containing the step.
User	Displays the run owner at the time the step was run.
Worker	Displays the UUID of the worker that executed the step.

Run Log > Find dialog box

Use the **Find** button in the Run Log to help locate a step that you want to see. Enter one or more search criteria, scrolling down to see those that are hidden. Steps that correspond to the search criteria are highlighted.

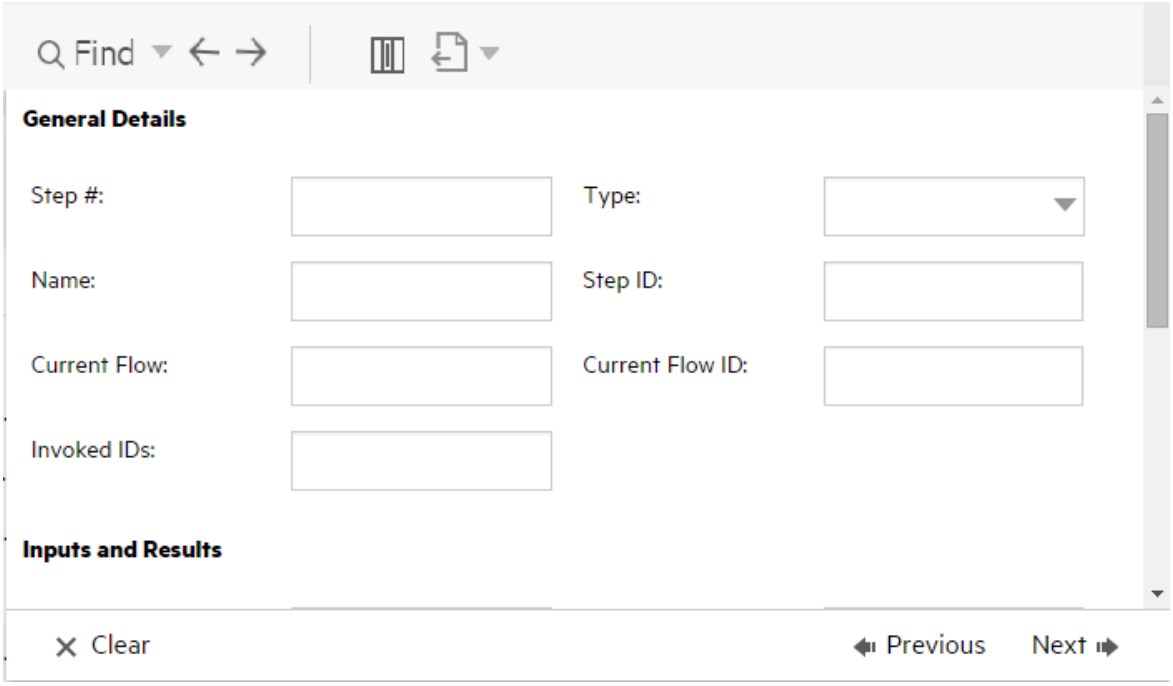

GUI item	Description
Clear	Click to clear the search criteria boxes.
Next	Click to highlight the next step that corresponds to the search criteria.
Previous	Click to highlight the previous step that corresponds to the search criteria.
← →	You can close the Find dialog box and use the arrows in the toolbar to navigate across the search results.

Flow Information

Click the **Down** arrow
▼ to the right of the flow run name, to expand the **Flow Info** pane.

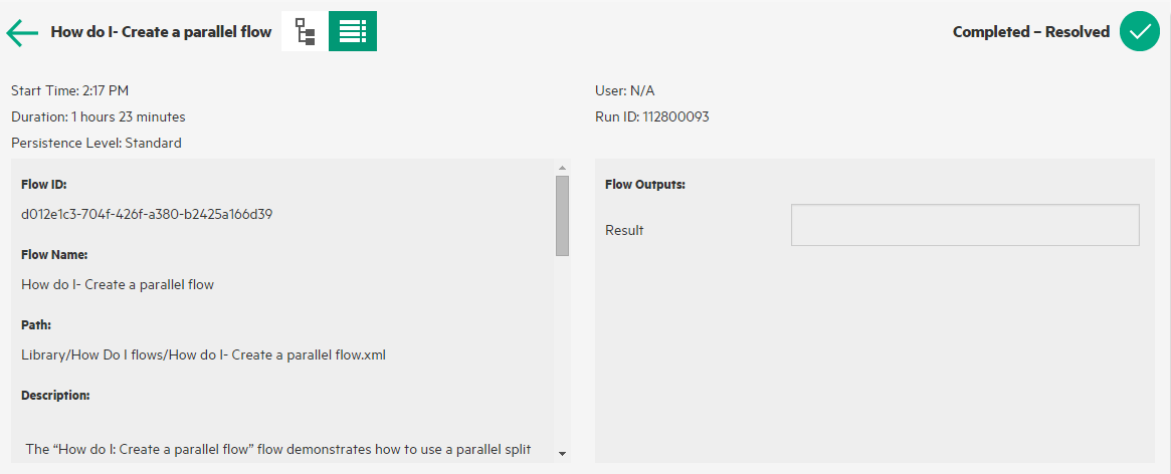

GUI item	Description
General details	Displays basic information about the run, such as start time, duration, user, run ID, and the persistence level, which specifies how the run history is saved to the database. For information about the persistence level, see "Monitoring and Controlling Database Size" on page 112.
Flow ID	Displays the UUID (unique identifier) of the flow.
Flow Name	Displays the name of the flow.
Path	Displays the path to where the flow is located.
Description	Displays the description of the flow, if one exists. If you can't see all the text in the pane, use the slider to display the text that is lower down in the pane.
Flow Inputs	Displays the flow inputs (if any exist).
Flow Outputs	Displays the flow output.
Up arrow	Click to collapse the **Flow Info** pane.

Embedding Central Views in External Web Pages

HPE OO gives you the option to embed parts of the Central UI into an external application. You can embed:

- Flow Run view - for running and interacting with flows
- Drill Down view - for advanced tracking of running flows

To embed the view in your web page, use an iFrame element in the HTML with the appropriate URL.

In the example below, the **999** in **drilldown-ex.html#999** stands for the **<run_id>** value. For more information, see Run a flow in an embedded Flow Run view.

```
<!DOCTYPE html>
<html lang="en-US">
<head>
<style>
 iframe {
 width : 800px;
 height : 600px;
 }
</style>
<title>HTML iframe tag</title>
</head>
<body>
<iframe src="http://localhost:8080/oo/drilldown-ex.html#999"></iframe>
</body>
</html>
```

Important! If you are going to embed the Drill Down window into an external HTML page, it is recommended to have an SSO solution in place (for example, SAML 2.0). For more information, see "Setting Up Security – LWSSO" on page 77 and "Setting Up Security – SAML Authentication" on page 70.

If an SSO solution has not been configured, an HPE OO login screen will appear in the iframe, when it first loads inside the page. On some browsers (for example, IE11), the default privacy settings will need to be updated to allow third party cookies before users will be able to log into HPE OO via the iframe.

What do you want to do?

Run a flow in an embedded Flow Run view

The Flow Run view lets you run a specific flow and perform minimal tracking as it runs.

After the flow is run, the basic tracking view appears.

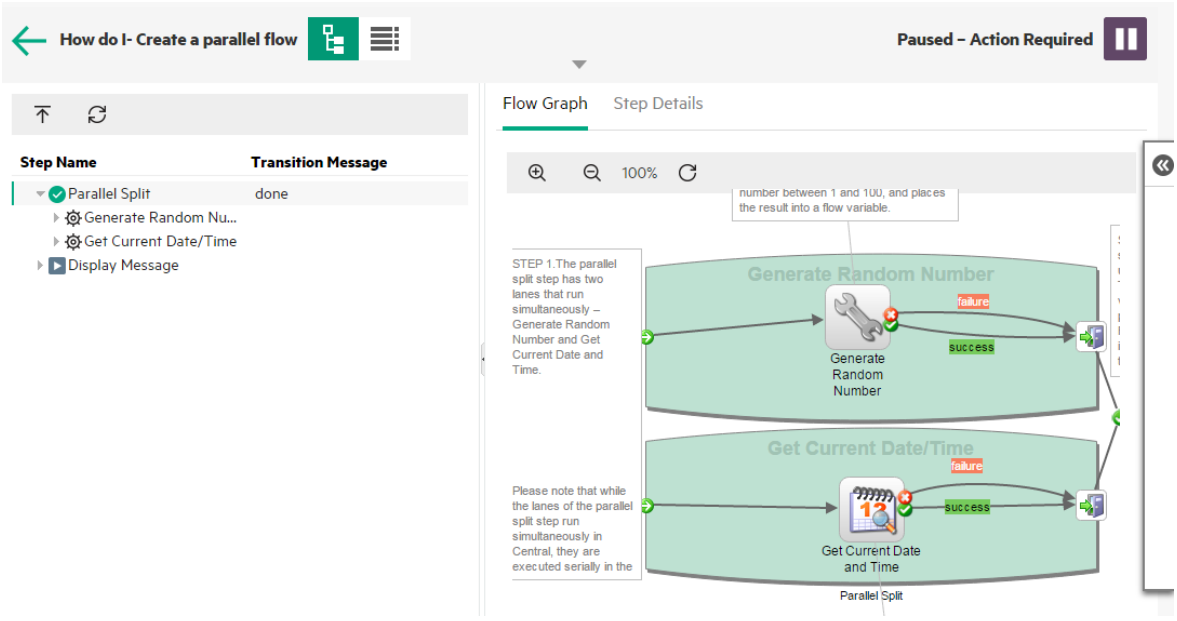

To run a flow in an embedded Flow Run view, use the following URL pattern:

```
http://<hos_name>:<port>/oo/trigger.html#<flow_uuid>
```

Identify the flow to run via the `<flow_uuid>`. You can find it in the flow's **Run Link** section in the Flow Library.

If required, use the `input_` and `runName` prefixes to edit the run link. You can get it to pass inputs or to change the name of the flow run.

The format is `input_<inputName>=<inputValue> and runName=<runNameValue>`. These values should come after the `?` and concatenate with `&`.

For example:

```
http://<hos_name>:<port>/oo/trigger.html#7a1fc3c7-1c5b-4868-a1c7-
2932d878897e?input_host=someHost&input_username=someUsername
```

Track a flow in an embedded Drill Down view

The embedded drill down view lets you track and interact with a running flow.

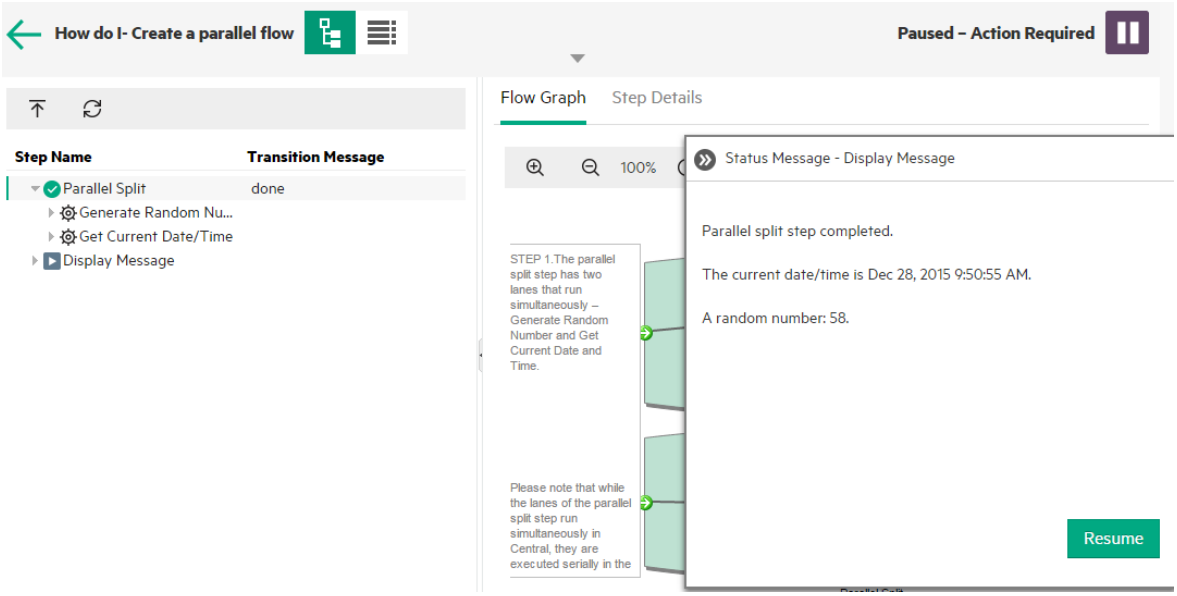

To track a flow in an embedded Drill Down view, use the following URL:

```
http://<hos_name>:<port>/oo/drilldown-ex.html#<run_id>
```

Identify the flow to track via the `<run_id>`. You can find it in the flow's **Run Link** section in the Flow Library.

Typical Central Errors

Typical Deployment Errors

Failed to read the file <cp jar name>. The file might be corrupt.

This error message appears when HPE OO is unable to read a content pack.

Solution:

If the file was copied from another machine, try copying it again.

Otherwise, go to Studio and select **Create content pack**, and redeploy the content pack.

Failed to read the file <cp jar name>. The file is empty.

This error message appears when the content pack file is empty.

Solution:

If the file was copied from another machine, try copying it again.

Otherwise, go to Studio and select **Create content pack**, and redeploy the content pack.

Wrong file format. <cp jar name> should contain the Lib and Content root folders.

This error message appears when a content pack is missing the required **Lib** and **Content** folders.

Solution:

If the file was copied from another machine, try copying it again.

Otherwise, go to Studio and select **Create content pack**, and redeploy the content pack.

> **Note:** Do not edit the contents of the content pack manually.

Wrong file format. <cp jar name> is missing the 'contentpack.properties' file.

This error message appears when a content pack is missing a required file.

Solution:

If the file was copied from another machine, try copying it again.

Otherwise, go to Studio and select **Create content pack**, and redeploy the content pack.

> **Note:** Do not edit the contents of the content pack manually.

Missing '<flow or operation/type of configuration item>' with UUID '<uuid>', which is required by the '<flow/operation>': '<flow path/operation uuid>'.

This error message appears when there is a missing flow, operation, or configuration item.

Solution:

1. Go to Studio and search for the flow, operation, or configuration item with the given UUID in either the HPE content packs or other projects in your repository.
2. Redeploy the deployment unit with the content pack that contains the flow, operation, or configuration item that you found.

Missing '<flow or operation/type of configuration item>' with UUID '<uuid>', which is required by the '<flow/operation>': '<flow path/operation uuid>' and <number of additional dependents> more . See the server log for the full list of dependents.

This error message appears when there is a missing entity (flow, operation, or configuration item) that other flows or operations depend upon.

Solution:

1. Check the server log to see all of the flows and operations that depend on this missing entity.
2. Go to Studio and search for a flow, operation, or configuration item with the given UUID in either the HPE content packs or other projects in your repository.
3. Redeploy the deployment unit with the content pack that contains the flow, operation, or configuration item that you found.

An exception occurred during deployment. Check the server log for more details.

This error message may appear in a number of different situations, leading to an exception during deployment. This includes conflicts between two deployed configuration items with the same type and path.

Solution:

Check the server log to see the details of the exception. In the case of a configuration item path conflict, check the **general.log** file for the most detailed information.

The content pack that you are deploying has been deployed previously. In the previous deployment, this content pack contained an entity ('<entity type>' named '<entity name>' with UUID <uuid>), which is used by other deployed content packs (<names of one or 2 cps>). The current content pack that is being deployed is missing that entity.

This error message appears when you are deploying a content pack that has been deployed previously, and this content pack originally contained an entity that is used by other deployed content packs. However, the version that you are deploying is missing that entity.

Solution:

Option 1: Go to Studio and search for the flow, operation, or configuration item with the given UUID in a previous version of the content pack that you are deploying, and add it to the current version of the content pack.

Option 2: Go to Studio and in the dependent content pack(s), remove the dependency on the given flow, operation, or configuration item, and redeploy that/those content packs as well.

The content pack that you are deploying has been deployed previously. In the previous deployment, this content pack contained an entity (('<entity type>' named '<entity name>' with UUID <uuid>), which is used by other deployed content packs (<name of one of the cps> and <number of other cps> more. The current content pack that is being deployed is missing that entity. See the server log for the full list of dependent content packs.

This error message appears when you are deploying a content pack that has been deployed previously, and this content pack originally contained an entity that is used by other deployed content packs. However, the version that you are deploying is missing that entity.

Solution:

Check the server log to see all of the content packs that depend on this missing entity.

Option 1: Go to Studio and search for the flow, operation, or configuration item with the given UUID in a previous version of the content pack that you are deploying, and add it to the current version of the content pack.

Option 2: Go to Studio and in the dependent content pack(s), remove the dependency on the given flow, operation, or configuration item, and redeploy that/those content packs as well.

A flow with UUID <uuid> has <number of scheduled tasks> scheduled tasks that will be deleted.

This warning message appears when a flow is being deleted by deploying a new version of an already deployed content pack that contained flows with scheduled tasks.

Solution:

There is nothing to fix.

Schema validation error in '<flow name>'. <the xml parser exception>.

This error message appears when the schema is not correctly validated.

Solution:

Go to Studio and select **Create content pack** and redeploy the content pack.

> **Note:** Do not edit the contents of the content pack manually.

A flow with a duplicated path is being deployed. This path already exists either in a flow in another content pack in the deployment unit, or in a flow that was deployed previously. To see the error details, run the deployment again with the server log level set to DEBUG.

This error message appears when you are deploying a content pack with a flow that has the same path as another flow in another content pack in the deployment unit.

Solution:

1. To see the details about the flow path, flow UUID, and content pack name, run the deployment again with the server log level set to DEBUG. You will receive one of the following errors:
   - A flow with a duplicated path: '<flow path>' is being deployed. A flow with this path has previously been deployed.
   - A flow with a duplicated path: '<flow path>' is being deployed. A flow with this path exists in the content pack '<name of content pack>'.
2. In Studio, modify one of the content packs so that there is no longer a duplication, and redeploy the content pack. For more information about the required changes, see the sections about these error messages, below.

A flow with a duplicated path: '<flow path>' is being deployed. A flow with this path has previously been deployed.

This error message appears when you are deploying a content pack with a flow that has the same path

as a flow that was previously deployed.

Solution:

In Studio, change the path of the flow (change the name of one of the folders or the flow name) and redeploy the content pack.

A flow with a duplicated path: '<flow path>' is being deployed. A flow with this path exists in the content pack '<name of content pack>'.

This error message appears when you are deploying a content pack with a flow that has the same path as another flow in another content pack in the deployment unit.

Solution:

In Studio, change the path of the flow (change the name of one of the folders or the flow name) and redeploy the content pack.

A '<flow or operation/type of configuration item>' with a duplicated UUID: <uuid> is being deployed. This UUID already exists in a '<flow or pperation/type of configuration item>'in the content pack '<name of content pack>'.

This error message appears when you are trying to deploy a content pack with a flow, operation, or configuration item that has the same UUID as an item in another content pack that is deployed in the deployment unit.

Solution:

1. In Studio, search for the flow, operation, or configuration item with the given UUID.
2. Copy the flow, operation, or configuration item into the same project.
3. Delete the original copy of the flow, operation, or configuration item, in that same project.
4. Select **Create content pack**, and redeploy the content pack.

> **Note:** Do not edit the contents of the content pack manually.

The property <content.pack.name/content.pack.version> is missing in the properties file contentpack.properties.

This error message appears when the content pack name or version are missing from the content pack's properties file.

Solution:

Go to Studio and select **Create content pack**, and redeploy the content pack.

Service unavailable. The server may be down, too busy, or experiencing network problems. Check the server log for the deployment status.

This error message appears when the server is either too busy to handle the request or is down for maintenance, or when some network error has prevented the response from arriving to the client.

Solution:

Check the server log to see if the deployment didn't start or end successfully, or failed.

During deployment the file name became corrupted. Please try to redeploy.

In some cases, the file name arrives in a corrupt state, and causes an exception.

Solution:

Deploy again.

Typical Central Errors

You do not have permission to perform this action.

This error message appears when you are logged in as a user that does not have permission to perform the action.

Solution:

Log in as a user with the appropriate permission. Alternatively, ask the administrator to give your user permission to perform the action.

An error has occurred. Check the server log for more details.

This error message appears when there is an unexpected error status from the server, but no error status text was received.

Solution:

Check the server logs to see if there was an exception.

liquibase: Waiting for changelog lock.

After a few attempts to log into the system, the server does not start and this error message appears.

Solution:

In the **DATABASECHANGELOGLOCK** table, set the value of the **LOCKED** column to **0**.

<server response status (for example, Internal Server Error)>. Check the server log for more details.

This error message appears when there is an unexpected error status from the server (for example, “Out of memory”).

Solution:

Check the server logs to see if there was an exception.

Hewlett Packard
Enterprise