

HP Operations Orchestration Software Central

Software Version: 7.60

Users' Guide

Document Release Date: January 2010

Software Release Date: January 2010



Legal Notices

Warranty

The only warranties for HP 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. HP 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 HP 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 Notices

© Copyright 2005-2010 Hewlett-Packard Development Company, L.P.

Trademark Notices

For information on open-source and third-party software acknowledgements, see in the documentation set for this release, Open-Source and Third-Party Software Acknowledgements (3rdPartyOpenNotices.pdf).

On the Web: Finding OO support and documentation

There are two Web sites where you can find support and documentation, including updates to OO Help systems, guides, and tutorials:

- The OO Support site
- BSA Essentials Network

Support

Documentation enhancements are a continual project at Hewlett-Packard Software. You can obtain or update the HP OO documentation set and tutorials at any time from the HP Software Product Manuals Web site. You will need an HP Passport to log in to the Web site.

To obtain HP OO documentation and tutorials

1. Go to the HP Software Product Manuals Web site (<http://support.openview.hp.com/selfsolve/manuals>).
2. Log in with your HP Passport user name and password.
OR

If you do not have an HP Passport, click **New users – please register** to create an HP Passport, then return to this page and log in.

If you need help getting an HP Passport, see your HP OO contact.

3. In the **Product** list box, scroll down to and select **Operations Orchestration**.
4. In the **Product Version** list, click the version of the manuals that you're interested in.
5. In the **Operating System** list, click the relevant operating system.
6. Click the **Search** button.
7. In the **Results** list, click the link for the file that you want.

BSA Essentials Network

For support information, including patches, troubleshooting aids, support contract management, product manuals and more, visit the following site: <http://www.hp.com/go/bsaessentialsnetwork>

This is the **BSA Essentials Network** Web page. To sign in:

1. Click **Login Now**.
2. On the **HP Passport sign-in** page, enter your HP Passport user ID and password and then click **Sign-in**.
3. If you do not already have an HP Passport account, do the following:
 - a. On the **HP Passport sign-in** page, click **New user registration**.
 - b. On the **HP Passport new user registration** page, *enter the required information and then click **Continue***.
 - c. On the confirmation page that opens, check your information and then click **Register**.
 - d. On the **Terms of Service** page, read the Terms of use and legal restrictions, select the **Agree** button, and then click **Submit**.
4. On the **BSA Essentials Network** page, click **Operations Orchestration Community**.

The Operations Orchestration Community page contains links to announcements, discussions, downloads, documentation, help, and support.

Note: Contact your OO contact if you have any difficulties with this process.

In OO: How to find Help, PDFs, and tutorials

The HP Operations Orchestration software (HP OO) documentation set is made up of the following:

- **Help for Central**
Central Help provides information to the following:
 - Finding and running flows
 - For HP OO administrators, configuring the functioning of HP OO
 - Generating and viewing the information available from the outcomes of flow runsThe Central Help system is also available as a PDF document in the HP OO home directory, in the \Central\docs subdirectory.
- **Help for Studio**
Studio Help instructs flow authors at varying levels of programming ability.
The Studio Help system is also available as a PDF document in the HP OO home directory, in the \Studio\docs subdirectory.
- **Animated tutorials for Central and Studio**
HP OO tutorials can each be completed in less than half an hour and provide basic instruction on the following:
 - In Central, finding, running, and viewing information from flows
 - In Studio, modifying flowsThe tutorials are available in the Central and Studio subdirectories of the HP OO home directory.
- **Self-documentation for operations and flows in the Accelerator Packs and ITIL folders**
Self-documentation is available in the descriptions of the operations and steps that are included in the flows.

Table of Contents

Warranty	ii
Restricted Rights Legend	ii
Trademark Notices	ii
On the Web: Finding OO support and documentation.....	iii
Support	iii
BSA Essentials Network.....	iii
In OO: How to find Help, PDFs, and tutorials.....	iv
This Help system and Guide.....	1
Quick View: Operations Orchestration Central.....	1
A scenario.....	2
Starting Central	2
Central Web application: the initial page	2
Flow Metrics area.....	3
Popular flows	4
Navigating in Central	5
Changing your OO Central password	5
Viewing the OO groups you belong to.....	6
Finding flows	6
Browsing flows in the Library.....	7
Searching for a flow	8

Previewing flows	9
Navigating in the flow preview	11
Three ways to run a flow	12
Running subflows	16
Opening the flow graph in a separate browser window	16
Seeing what has happened in the flow run	18
Run histories: What happened and why	21
Run histories in more detail	25
Viewing step results and other advanced data	27
Scheduling flows	28
Run-scheduling concurrency	29
Creating a schedule for a flow	29
Working with flow schedules	32
Editing existing schedules	33
Enabling and disabling existing schedules	33
Deleting schedules	33
Configuring Scheduler settings	33
The Dashboard: Learning more from flows	35
Collating data on Dashboard reporting charts	35
What do the bars tell you?	37
Learning more from the charts	38
Creating and modifying charts	39
Exporting and importing chart definitions	42

Starting a flow from outside Central	43
Interrupting flow runs	44
Reassigning ownership of a run.....	44
Resuming a run	45
Deleting a run.....	45
Creating a link to a flow run	46
Creating a link to a flow	47
Handing off flow runs.....	48
Auditing and managing flows.....	49
Users, groups, and access control	49
Capabilities and access permissions	50
Capabilities.....	50
Permissions.....	51
Allowing external users into the Central system	52
Using external authentication for Central users.....	53
AD authentication settings	53
LDAP authentication settings.....	57
Kerberos authentication settings.....	59
Managing users.....	61
Internal or external users?.....	61
Adding a user	62
Editing a user's account	63
Deleting a user	63
Managing groups	64
Default groups on the Manage Groups subtab	65
Adding groups.....	66
Adding OO users to groups.....	67
Mapping external groups to OO groups	68
Changing a group's assigned capabilities and description.....	69
Deleting groups.....	70
Managing flow runs.....	71
Other system configurations	72
Enabling the visibility of ROI reporting	73

Restricting authoring to administrative users.....	74
Setting how many versions of an object are stored	74
Restricting who can publish	75
Changing the Dashboard charts refresh rate	75
Specifying a required prefix for init params when using a URL to start a flow run in Central	75
Specifying how Central manages LDAP referrals	76
Enabling Central to resume runs that were interrupted when Central failed	76
Changing the configuration of a Central cluster	77

Troubleshooting.....77

Web browser shows a security-certificate-related warning when you open the Central Web site	77
I was sent back to the login page.....	77
I cannot restart or resume a flow that I started or that was handed off to me.....	78
I cannot change or create a schedule for a flow	78
Central fails to create a schedule for a flow	78
Changes were made to the Public repository, but they don't appear in the Flow Library.....	78
Flows run under heavy load with command-line or RAS operations fail on Windows systems with error code 128	78

Index79

This Help system and Guide

Help for Central (reproduced in the PDF *Operations Orchestration software Central Users' Guide*, Central_UsersGuide.pdf) provides an introduction to Central and detailed procedures that you will use to create flows.

This Help system is intended for all Central users. It provides a high-level overview of HP Operations Orchestration software (HP OO) and flows and detailed instructions on using Central. After reading the introduction, users can break out to those of the following chapters that are appropriate to what they will be doing and which component they will be using:

- Introduction to HP OO
This section is for all users; it gives an overview of HP OO and its concepts.
- Using Central
This section is for IT staff that run flows.
- Administering HP OO
This section covers administrative tasks.
- Viewing flow Reports and Audit Trails
This section is for IT managers who want to study metrics and reports on flows and runs.

Quick View: Operations Orchestration Central

HP OO Central provides automated sequences of tasks called *flows* that you run to reduce the time required to keep your organization's network functioning.

You access and run flows from Central to:

- Diagnose and repair network problems.
- Monitor the health of applications and networks.
- Perform maintenance tasks.

This Quick View of Central will show you how Central:

- Enables front-line IT support personnel to resolve alerts and repair tickets, check the health of applications, servers, and peripherals, and perform repeated maintenance tasks more quickly and with full auditing.

You can accomplish these goals with the flows in the Central Library. A flow is an automated, structured sequence of operations that can respond to the conditions it finds.

- Helps IT managers understand precisely where their system needs help and how the flows are doing at providing that help.

Dashboard reporting charts graphically relate incidents to the causes of problems. For example, you can chart which servers are going down more often than is normal. To learn what the underlying problem is and how it was solved, you can then look at the run histories for the flow that brought the server back up. Some services may be restarted many times a day without being logged anywhere. This information is now available with the OO charts and reports.

Further, you can drill down into the information that Central has recorded. For example you could examine your most common alerts to see which operating system they occur most frequently on, and then drill down further to see which particular system is most problematic.

Reporting charts and run histories also tell you whether a given flow is accomplishing what it's intended to do, or whether the flow author needs to work on it more.

A scenario

Suppose your IT department encounters a broad range of alerts that originate from various servers, applications, and operating systems. In addition to resolving the alerts, you need to mine meaningful data from the information that comes out of using various actions to resolve those alerts.

To see what you can do with Central, let's look at both of those goals:

- Central users run the flows that resolve the alerts.
- Users then analyze the data that is produced by the flows that resolve the alerts to discover information such as:
 - What are the alerts that are showing up most frequently?
 - What is the outcome for each alert?
 - Which server or application generated the most alerts?
 - Which flows ran most often, and what were their outcomes?
 - Which applications and servers had fatal errors?
 - How many alerts of various severities were there?

We'll look at both these goals in turn.

Starting Central

If you run Central in Internet Explorer on a machine running a Windows Server operating system, you must add the domain address of Central (<https://<your-hostname>>) to the Intranet Security Zone, using the default settings.

Central provides a graphical user interface for:

- Finding and running flows (on the **Flow Library** tab).
- Creating reports and viewing information on flow runs (on the **Reports** tab).
- Viewing and customizing charts of aggregated data that show patterns in flow usage and outcomes and may reflect the state of your IT system's health.
- Scheduling flow runs (on the **Scheduler** tab).
- Changing Central and OO configurations and enabling external authentication (on the **Administration** tab).

To start Central

1. Start your machine's Web browser.
2. Paste the URL that your administrator sent you into the **Address** box of the Web browser and then press Enter.
3. When the message appears that you are about to view pages over a secure connection, click **OK**. If a message appears, warning you that the site is not trusted, it is safe to proceed.
4. Click **Yes**.
5. When the Central **Login** page appears, log in with your user name and password. Central opens and you're ready to locate, run, and view information on flows.

Central Web application: the initial page

When you start the Central Web application, the default start page is the Dashboard, where you can analyze results of flow runs.

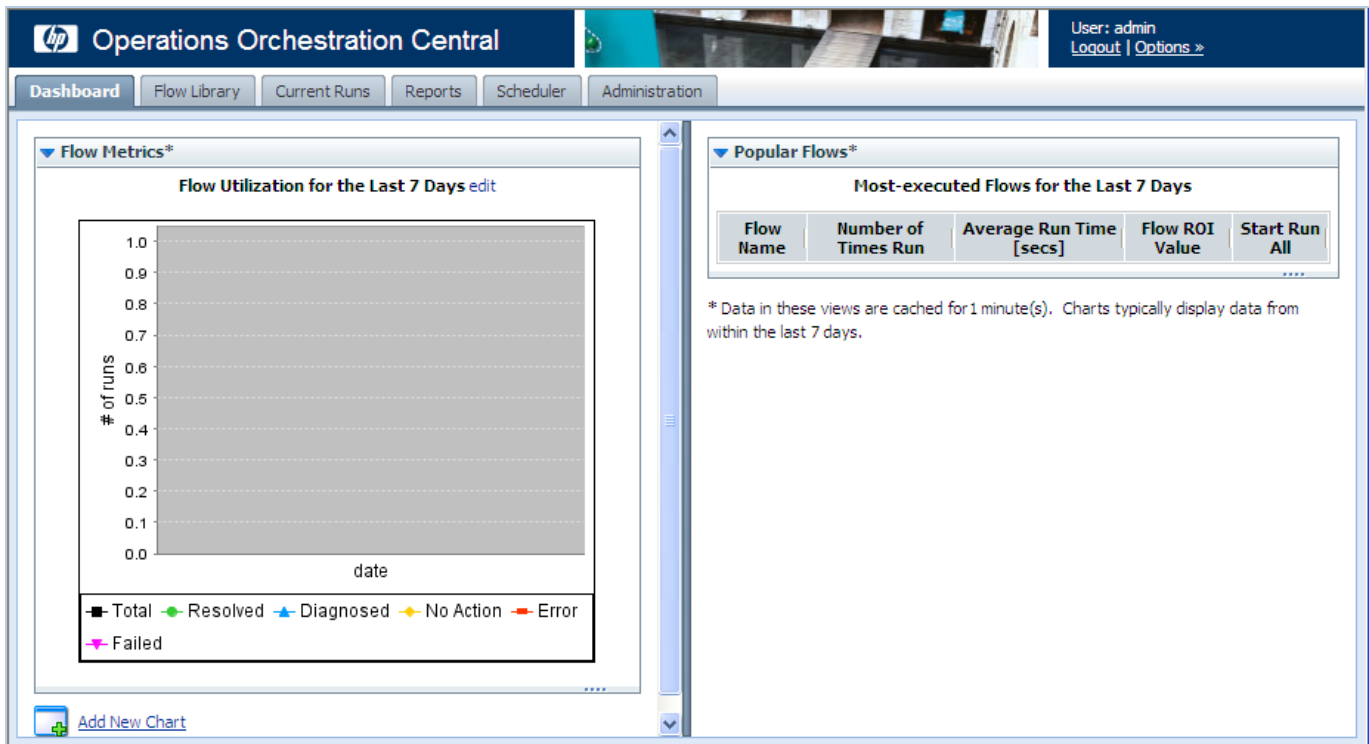


Figure 1 - Central Web application Dashboard

Note: Whether you see the **Current Runs**, **Reports**, or **Scheduler** navigation tabs depends on which capabilities (abilities to schedule flows, to manage other users' runs, etc.) are assigned to your group or groups.

The **Flow Metrics** area is a diagnostic and analytical area where you can call up and create charts that offer different views of information obtained by all the flows that have run. The **Popular Flows** area is where you can examine histories of flow runs.

Flow Metrics area

The **flow Metrics** graph shows one of the following metrics over the last week, month, or year:

- Total number of flow runs, broken down by the outcome of the run (problem resolved, problem diagnosed, no action necessary, error, or run failure).
- The average execution time of all the runs (MTTR, or Mean Time to Resolution).
- The total value of the run, as determined by the monetary value that the flow author assigned to completion of each step in the runs.

By default, it shows flow utilization (the number of flow runs) over the last 7 days, but you can change either what it shows or the time span that it covers.




Figure 2 - Flow Metrics area

To customize the flow Metrics graph

1. To change the information that the **Flow Metrics** graph displays or the time span that it reports, click **edit**.

The **Flow Metrics** editing area appears, in which you can choose the metrics that you want to see and the time span that you want the metrics to cover.



2. Select your choices in the list boxes beneath the graph.
If you select years, the intervals represented are months.
3. To update the graph, click **Go**.

Popular flows

The **Popular Flows** area provides a quick view of the executed flows that have recently been run the most, including a shortcut for running a flow again.

▼ Popular Flows*

Most-executed Flows for the Last 7 Days

Flow Name	Number of Times Run	Average Run Time [secs]	Flow ROI Value	Start Run All
Windows Health Check	1	43.297	.00	
Simple SMTP Check	1	16.172	.00	
Restart Service - Tutorial Flow	1	19.749	2.15	

*Data in these views are cached for 1 minute(s). Charts typically display data from within the last 7 days.

Figure 3 - Popular Flows

In addition to being able to start any of the flows listed here (by clicking the green arrow), you can also open any of the Dashboard charts.

To open a Dashboard chart

1. On the **Dashboard** tab, click **Add New Chart**.
2. In the **Select a report to view** drop-down list, select the chart that you want to see, and then click **View**.

Navigating in Central

The Central interface varies according to whether you are finding or running a flow or generating a report or metrics. However, you can always navigate with the **Dashboard, Flow Library, Current Runs, Reports, Scheduler, and Administration** tabs.




Figure 4 - Navigation tabs

Which navigation tabs you see depends on which capabilities have been assigned to the group or groups that your account is a member of. The ADMINISTRATOR group can see all the tabs. Otherwise:

- The **Dashboard** and **Reports** tabs appear only if your account’s group has the RUN_REPORTS capability.
- The **Scheduler** tab appears only if your account’s group has the SCHEDULE or VIEW_SCHEDULES capability.


Note: If your group has only the VIEW_SCHEDULES capability, you can see existing flow schedules on the tab, but they are read-only. To create or change schedules, your group must have the SCHEDUL capability.

- Everyone can see the Administration tab, but what you can see and do on it depends on your group’s capabilities.

Changing your OO Central password


You change your account’s password on the **Administration** tab.

To change the password for the current user

1. Click the **Administration** tab, and then click **Account** ( Account).
2. In the **Current Password** text box, type your current password.
3. In the **New Password** and **Confirm New Password** text boxes, type the new password, and then click **Change**.

Viewing the OO groups you belong to

To view the group memberships of the current user

- Click the **Administration** tab, and then click **Account** ( Account).
The **My Groups** panel lists the groups to which the logged-in user account belongs.

Finding flows

Your first question is probably which flow to run. You can either browse Central's Library on the **Flow Library** tab or use the **Search** feature to find the flow you needed to resolve each alert.

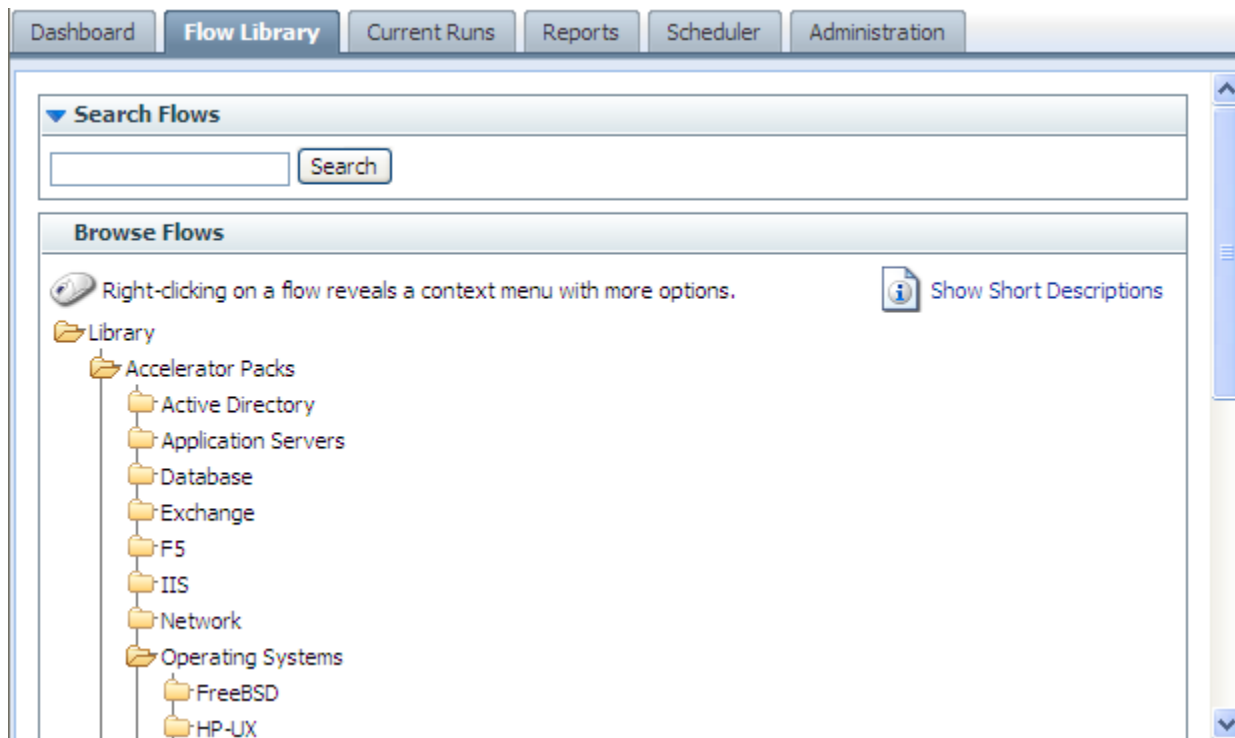



Figure 5 - The OO Library

Some of the folders in the Library group the flows according to the technology area in which they solve problems. The flows that come with your initial installation are organized this way, under the **Accelerator Packs** and **ITIL** folders. For instance, if you want to check your Internet Information Services (IIS) SMTP server health, you would expand the **Accelerator Packs** folder, the **IIS** folder, the **Utility** folder, and then run the **SMTP Server Health** flow.


If you can't find the flow you need, you can search for it.

Browsing flows in the Library

To browse the Library for a flow

1. Click the **Flow Library** tab.
The flow **Library** opens.
2. To find a flow, open the **Library** folder and navigate through the folder tree to the flow.
The  icon represents a flow.



Tip: To see short descriptions of what each flow does, click **Show Short Descriptions**. When the descriptions are displayed, the command changes to **Hide Short Descriptions**. To see more detailed information on the flow, click the “i” balloon ().

An information box such as the following appears, containing descriptions and other information about the flow.

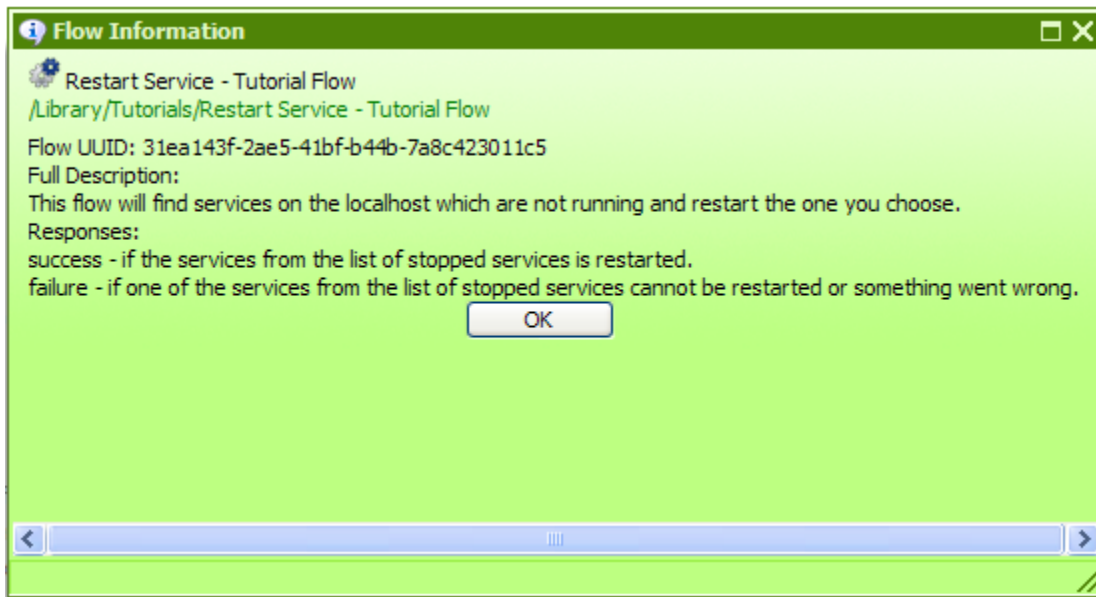


Figure 6 - More information on the Restart Service - Tutorial flow

Click **OK** to close the information box.

3. To run the flow, click the flow name.
This loads a preview of the flow and allows you to choose how you want to run the flow.
There are three ways to run a flow:
 - **Guided Run**
You click to carry out each step and respond to any user prompts.
 - **Run All**
The flow completes all the steps on its own, and you only respond to user prompts.
 - **Instant Run**
Like Run All, except that the user prompts, results, and other data generated appear in a dialog box.

For more on running a flow, see [Three ways to run a flow](#).

Searching for a flow

Central's search mechanism uses the Apache Lucene search syntax. In addition to the basic search method described in this topic's procedure, you can use the syntax to construct more highly targeted searches. For more information on the search syntax, see the Apache Software Foundation Web site's page on query syntax (<http://lucene.apache.org/java/docs/queryparsersyntax.html>).

To search for a flow

1. On the **Flow Library** tab, in the **Search flows** text box, type one of the following.

- Some or all of the flow's name
- Keywords
- A word or phrase within the flow description
- A flow category

OR

Type a search field and value using the form

`<fieldname>:<term>`

where `<term>` is the particular value in the field that may find the desired flow.

The fields that are available for searching are the following (they are not case-sensitive):

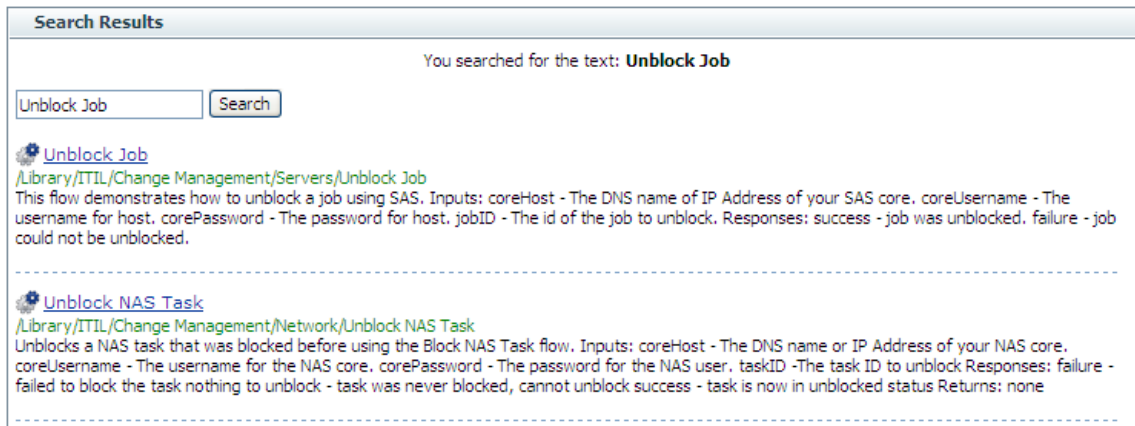
- **Category**
The category that has been assigned to the flow.
- **Description**
The flow's description.
- **Domain**
A domain term that has been associated with the flow.
- **ID**
The flow's Universally Unique ID.
- **Input**
An input to an operation used in the flow.
- **Name**
The flow's name.
- **Type**
The type of an operation used in the flow. The terms that you can match in this field and the operation types that they represent include:
 - cmd – Command
 - flow – An operation that is a flow
 - http – Http (also known as shell)
 - other - Scriptlet
 - script.perl – Perl script
 - ssh – SSH (Secure Shell)
 - telnet - Telnet
 - lock = Acquire Lock
 - unlock = Release Lock
- **Stepdescription**
The description of one of the flow's steps.

- **Stepname**

The name of one of the flow's steps.

2. Click **Search**.

The search results appear.



Note that the search results include:

- The description and inputs for each flow.
- Where the flow is located in the Library.

3. To load the flow from the results into the preview, click the flow name.

Previewing flows

The flow preview page contains the flow graph and information about it. When you run the flow step by step, the diagram illustrates the current progress of the flow.

The preview of the flow provides:

- A graph of the flow
- Buttons for starting the flow in various modes
- Various navigation icons.
- Panels with specialized information:
 - **Flow Details** panel: the flow's location in the Flow Library, its universally unique ID (UUID), and its description
 - **Reporting** panel: charts built from data reported by flow inputs
 - **Execution Links** panel: links that you can send to another program or user to start each of the kinds of flow runs
 - **Graph** panel: graphic depiction of the flow

To preview a flow

- In the folder tree on the **Flow Library** tab, click the flow.
The flow preview is loaded and appears.

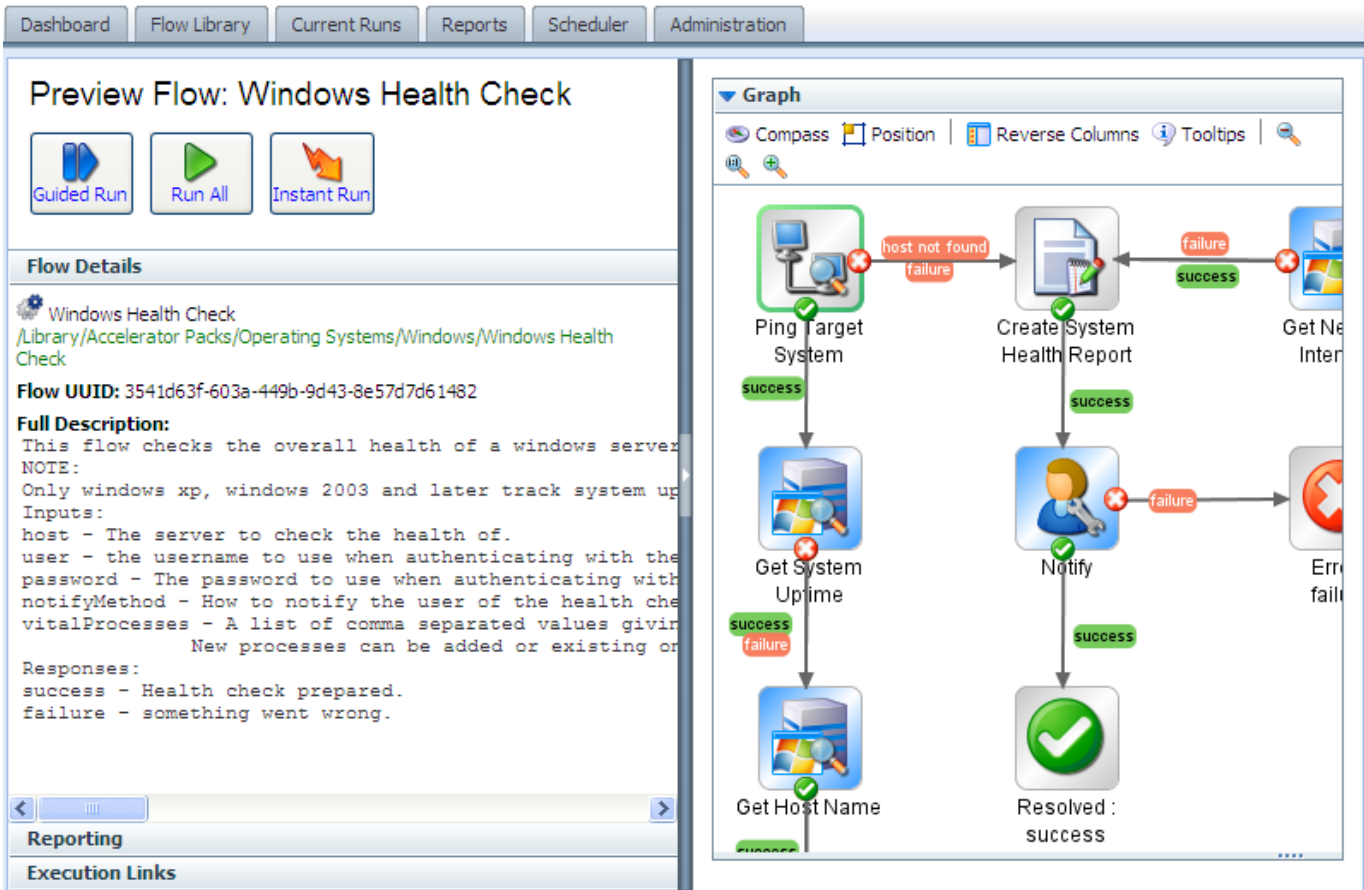


Figure 7 - A flow loaded and ready to run





Tips:






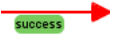
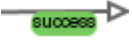
- To return to the Library without running the flow that you are previewing, just click the Flow Library tab again.
- To run the flow repeatedly, you can create a schedule for it. For information on scheduling repeated runs of a flow, see *Scheduling flows*.

The following table describes the symbols that may appear in the flow graph and the flow preview page.

In the graph of the flow, the following symbols have the meanings as described in this table:

Table 1. Flow graph symbols and their meaning

Symbol	Name	Meaning and comments
	Start step	The entry point for a flow. The green outline signifies that the step is the start step. (The flow operation symbol inside the step varies with the operation from which the step was made.)
	Parallel split step	A parallel split step is kind of hard to miss, because the lanes that follow tend to be rather large. Nonetheless, this symbol represents the start of the parallel split step. Parallel split steps are used to run distinct sequences of steps simultaneously.

Symbol	Name	Meaning and comments
	Multi-instance step	A multi-instance step performs the same action with different values for a given input (often specifying different targets), at the same time. For a multi-instance step, all the responses except one must ultimately loop back to the multi-instance step. The only response that leads ultimately to a return step is the group done response.
	Diagnosed return step	The step that ends a flow when a problem has been diagnosed.
	Error return step	The step that ends a flow when an error has occurred that prevents the flow from continuing. An error return step can also be used to indicate an expected condition is false (for instance if a counter is greater than a given number).
	Resolved return step	The step that ends a flow when a problem has been resolved.
	No action taken return step	The step that ends a flow when no action needs to be taken.
	Gated transition	A transition that is gated, or restricted to users with certain access permissions, appears in red on the canvas.
	Handoff transition	A transition after which the flow pauses for handoff to another user.


The other icons in the graph enable you to move around in the graph. For more information, see [Navigating in the flow preview](#).

Navigating in the flow preview

To show more of the flow graph

- Drag the vertical resizing bar in the middle of the page.

To zoom in or out of the flow graph

- In the bar at the top of the diagram, click the **Zoom In**, **Restore**, or **Zoom Out** icon () , according to the view you want.

To move the flow graph on the page


1. Click the **Compass** icon ( **Compass**) in the top-left of the flow graph. The compass appears.



Figure 8 - The compass

2. To move the flow graph, click the directional buttons in the compass.
3. To return the flow to its original position, click the center of the compass.

To reverse the panels on the preview page

- Click the **Reverse Columns** button ( Reverse Columns).

To give the flow graph top-left orientation

- Click the **Position** button ( Position).


To see descriptions of each step

- Click the **Tooltips** button ( Tooltips).

As you move the cursor over a step, its description appears.

Three ways to run a flow


From within the Central Web application, there are three kinds of flow runs and several ways to start each kind of run:

- **Guided run:** 

In a guided run, you click to carry out each step.

- **Run all:** 

In run all, when you start the flow, it runs straight through to completion, except for any user-prompt inputs that require you to supply a value.

- **Instant run:** 

An instant run is useful when you want to start a run without leaving the Flow Library or preview page. An instant run opens a dialog box, such as the following one for the **Windows Health Check** flow, within which the new run starts and runs to completion.

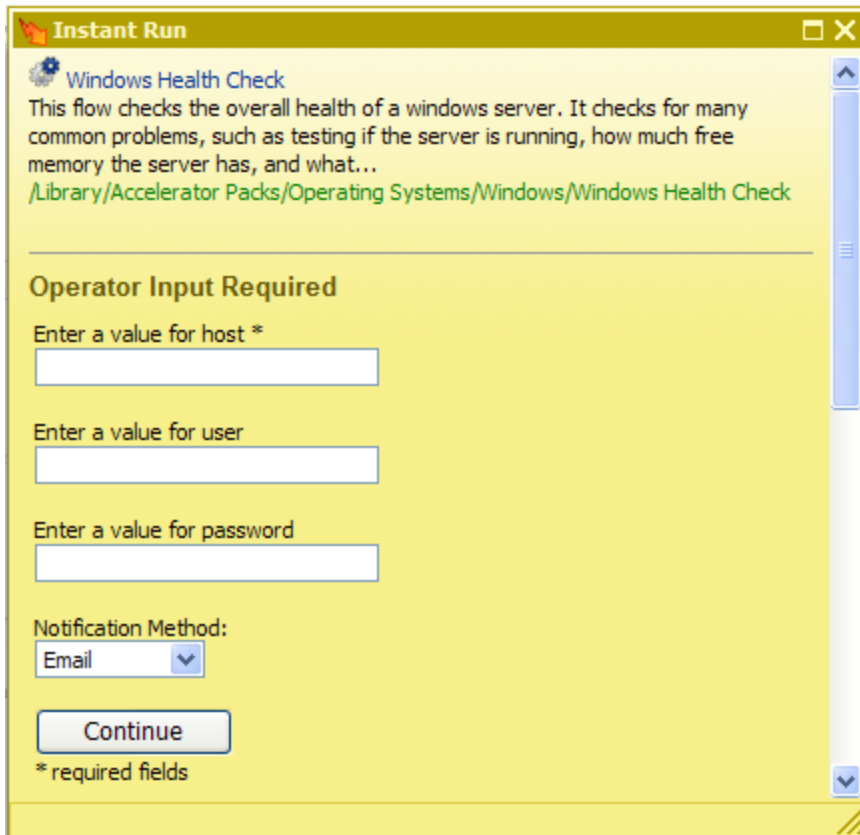


Figure 9 - Instant run of a flow

Only the essential controls for completing the run are available, and only while the box is open. Note though, the information that is visible when you expand the **Instant Run** dialog box.

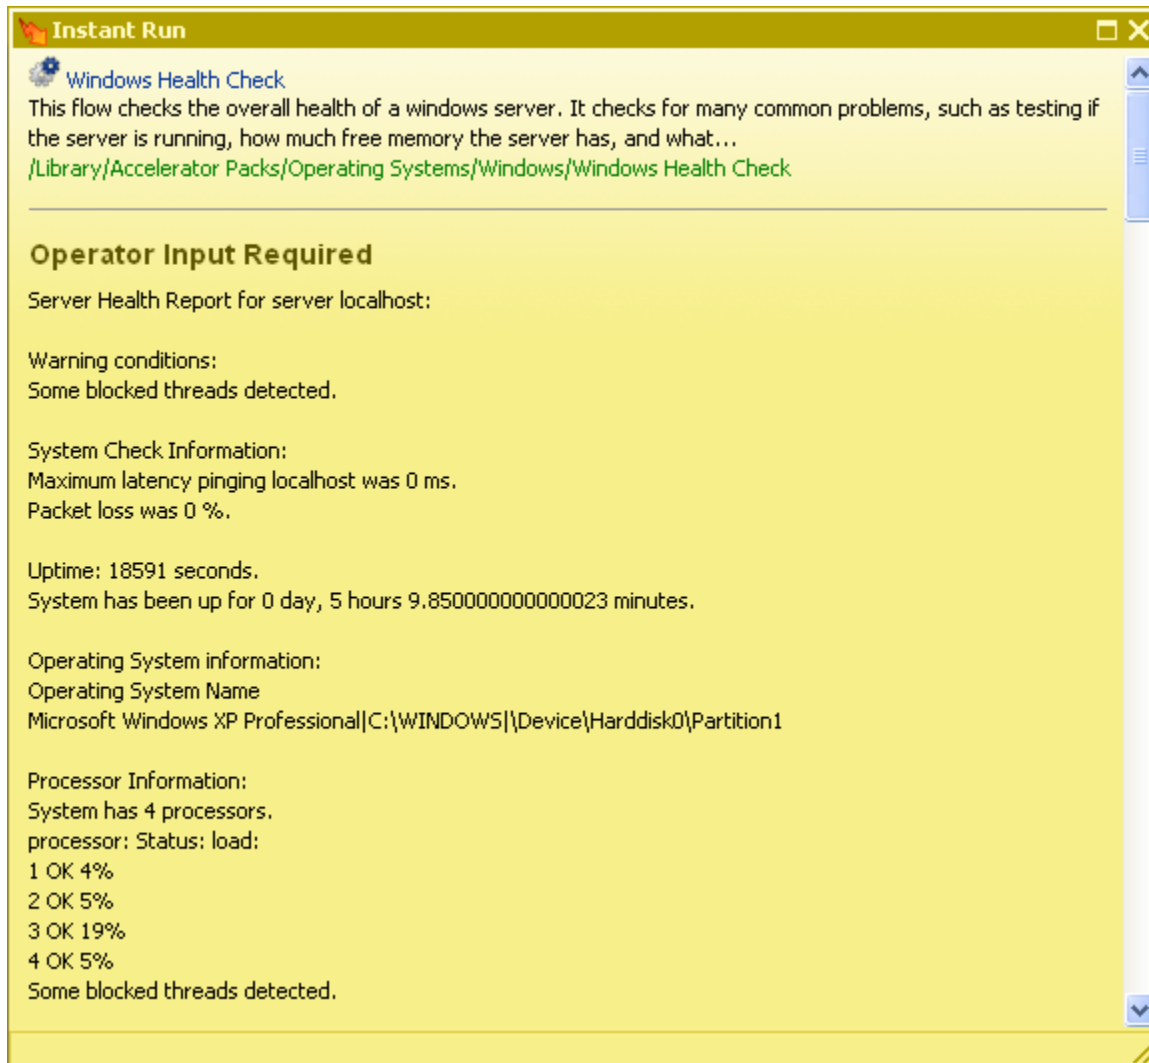
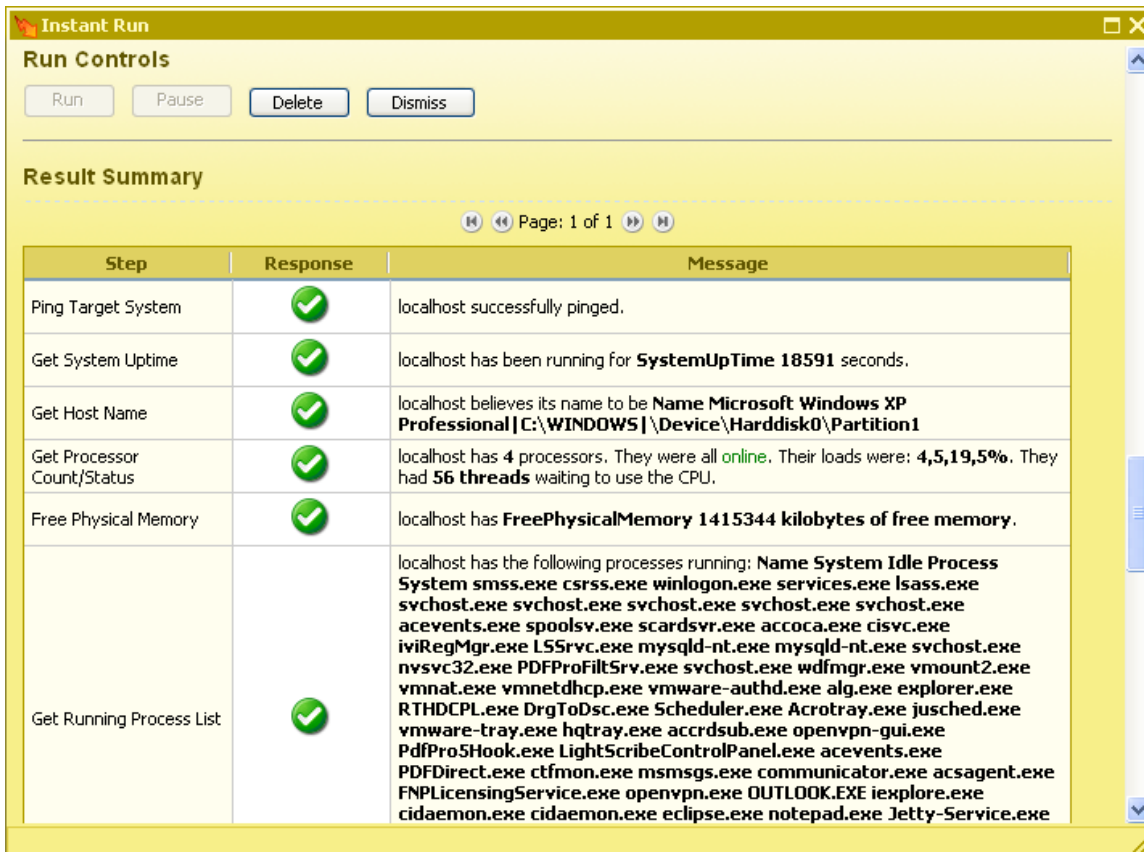


Figure 10 - Instant Run dialog box, expanded

Note: If your account does not have sufficient permissions or the capability to start a run of this flow, a message appears, alerting you do the fact and providing a button for handing off the flow. If you need to provide credentials, the username and password text boxes and Login button appear in the Instant Run dialog box.

As you scroll down, you'll find more information, such as the results for steps where relevant and the summary information for each step.



You can also start a flow run and search for flows from outside the Central application. For more information, see [Starting a flow from outside Central](#).

To run a flow

1. On the **Flow Library** tab, navigate to the flow you want to run.
2. Click the flow name to open the flow's Preview.
OR
Right-click the flow name.
3. On the preview page or (on the **Flow Library** tab) from the right-click menu, click one of the following:
 - **Guided Run**
 - **Run All**
 - **Instant Run**

If you choose **Guided Run** or **Run All**, the flow begins execution and the page contains the following panels:

- **Results Summary** panel: shows the step names, their responses, and what happened during each step
- **History Tree** panel: shows the order in which the flow steps and subflow steps were run and the results of the steps
- **Context** panel: shows all of the flow variables used in the flow run and the values assigned to them
- **Stack** panel: shows the root flow and, if it uses a nested-execution pattern, which subflow you may currently be inside
- **Graph** panel: graphic depiction of the flow that highlights each step as it is executing

If you choose **Instant Run**, the **Results Summary**, **History Tree**, **Context**, and **Stack** panels will not appear on the page, and the graph will not highlight the steps as they run.

Running subflows

All steps are based on operations, and a step's operation can itself be a flow. The flow that is the basis for the step is then considered to be a subflow of the flow that contains the step. The flow that contains the step is considered to be the parent flow.

When you are running a flow step-by-step (in a guided run) and come to a step that contains a subflow, you can step into the subflow or run it as a single step.

To step into and out of a subflow

1. When the step highlight moves to the step that contains the subflow, click **More Controls**.



The **Step Into** icon appears among the other additional controls.



2. Click **Step Into**.
3. Complete the steps of the subflow using the same procedure as you do for completing the steps of the parent flow.

The **Step Out** icon joins the others.



4. To run the subflow to completion and return to the steps of the parent flow, click **Step Out**.
OR

If you have reached the end of the subflow, click **Next Step**.

5. To complete the run, continue completing steps as described above.

Note: At any time, you can run the rest of the flow by clicking **Run All**.

Opening the flow graph in a separate browser window

When you run a flow step-by-step, you can give yourself more room for looking both at the data that the flow generates and that appears under **Results Summary** and at the flow graph by opening the flow graph in a separate window. The larger the flow, the more data that it generates, the more useful this can be. Consider the following example.

Dashboard | Flow Library | Current Runs | Reports | Scheduler | Administration

Run Status: STOPPED Started: 10:56:33 PDT 2009-06-18
 Current Step: Get System Uptime Options

more controls »

Results Summary

Last Completed Step: **Ping Target System** As of 10:58:05 PDT 2009-06-18

Page: 1 of 1

Step	Response	Message
Ping Target System		localhost successfully pinged.

Display most recent: ON 50 Items per page

History Tree

- Context
- Stack

Flow

Windows Health Check
 This flow checks the overall health of a windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what...
 /Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check

Graph

Compass | Position | Open Graph | Tooltips

```

  graph TD
    A[Get System Uptime] -- success --> B[Ping Target System]
    B -- success --> C[Create System Health Report]
    C -- success --> D[Notify]
    D -- failure --> E[Get System Uptime]
  
```

Figure 11 - Flow graph in same window

Here is the same flow after you click **Open Graph**.

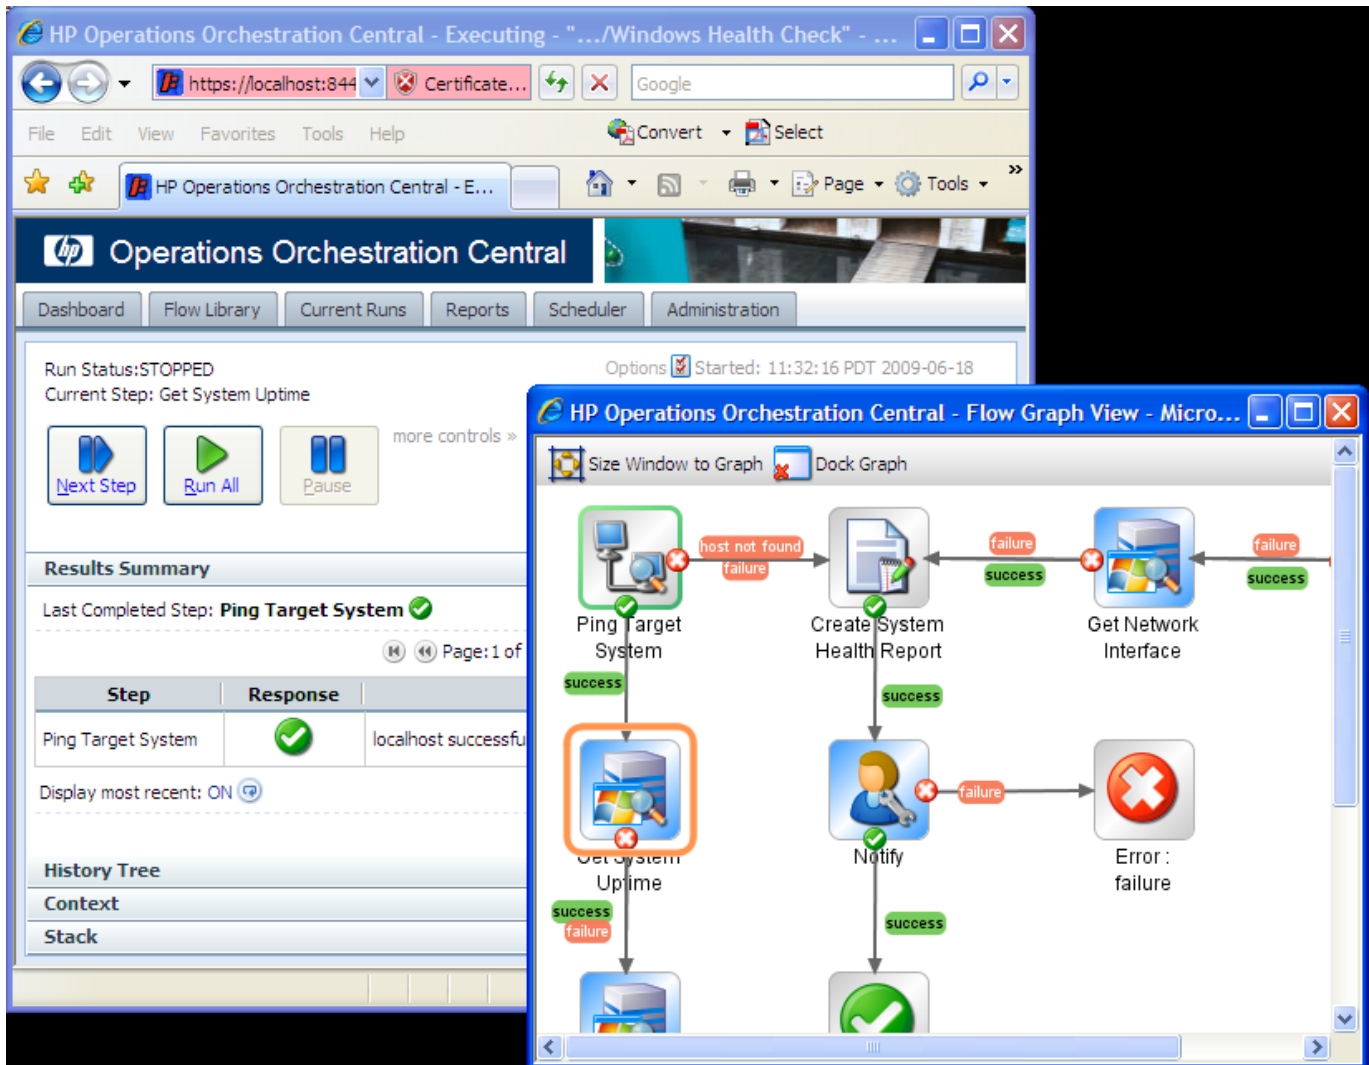


Figure 12 - Flow graph undocked from Library tab

Providing even a small flow with its own browser window provides much more room for the **Results Summary** and other information provided on the flow run.

To open the flow graph in a different browser window

1. With the flow running in Guided Run mode, above the flow graph, click **Open Graph**.
2. To size the new browser window to fit the flow graph, click **Size Window to Graph**.
3. To return the graph to the original Central browser, click **Dock Graph**.

Seeing what has happened in the flow run

While a flow is running, you can quickly see what has happened so far in the run.

The **Results Summary** panel, below, shows the step name, the response, and what happened in the step, as reported by the description in the transition that followed the step.

Run Status: READY Started: 15:43:40 PST 2009-02-02
 Current Step: Get Host Name Options

Next Step Run All Pause more controls »

Results Summary

Last Completed Step: **Get System Uptime** As of 15:46:29 PST 2009-02-02

Page: 1 of 1

Step	Response	Message
Ping Target System	✓	localhost successfully pinged.
Get System Uptime	✓	localhost has been running for SystemUpTime 17266 seconds.

Display most recent: ON 50 Items per page Update

History Tree Context Stack

Flow

Windows Health Check
 This flow checks the overall health of a windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what...

Graph

Compass Position Open Graph Tooltips

The flow graph illustrates the execution of a 'Windows Health Check' flow. It starts with 'Ping Target System' (success), followed by 'Create System Health Report' (success). From there, it branches to 'Get System Uptime' (success) and 'Notify' (success). A 'failure' path leads from 'Notify' to 'Get N...' (failure). Another 'failure' path leads from 'Ping Target System' to 'Create System Health Report' (labeled 'host not found failure').

Figure 13 - Results Summary panel, mid-run

Following is a close-up of the Results Summary.

Results Summary

Finished Run As of 15:55:25 PST 2008-11-20

Page: 1 of 1

Step	Response	Message
Get Stopped Services	✓	Retrieved a list of services which are currently stopped.
Select a Service	✓	Selected .NET Runtime Optimization Service v2.0.50727_X86 to restart.
Restart Service	✓	Restarted service .NET Runtime Optimization Service v2.0.50727_X86
Resolved : success	✓	Return step - Restart Service - Tutorial Flow

Display most recent: ON 50 Items per page Update

History Tree Context Stack

You can also view a flow's run history on the **Reports** tab.

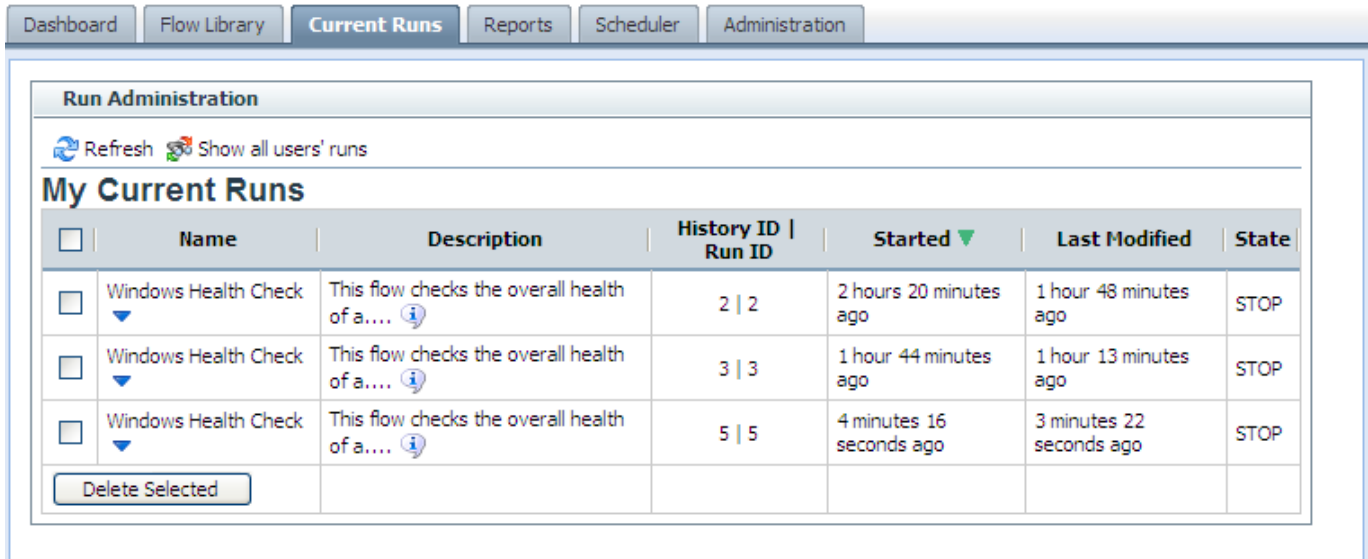
Necessary capabilities:

- Your group must have the RUN_REPORTS capability to view the **Reports** tab.

- To view runs other than your own on the **Current Runs** tab, your group must have the `MANAGE_RUNS` capability.

To view a flow's run history mid-run

- Click the **Current Runs** tab.



To delete a flow run mid-run

- On the **Current Runs** tab, select the box for the run that you want to delete.
- Click **Delete Selected**.

To view the run's history mid-run

- On the **Current Runs** tab, click the downward-pointing arrow for the run that you want to view. A drop-down menu appears, including the command **Inspect History**.

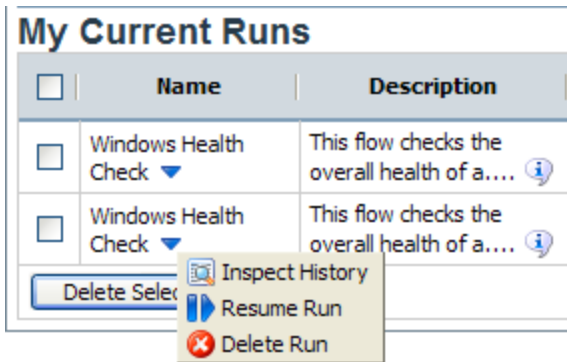


Figure 14 - Runs in the Current Runs tab

- Click **Inspect History**.
The run history is a report for the single run, which opens on the **Reports** tab.

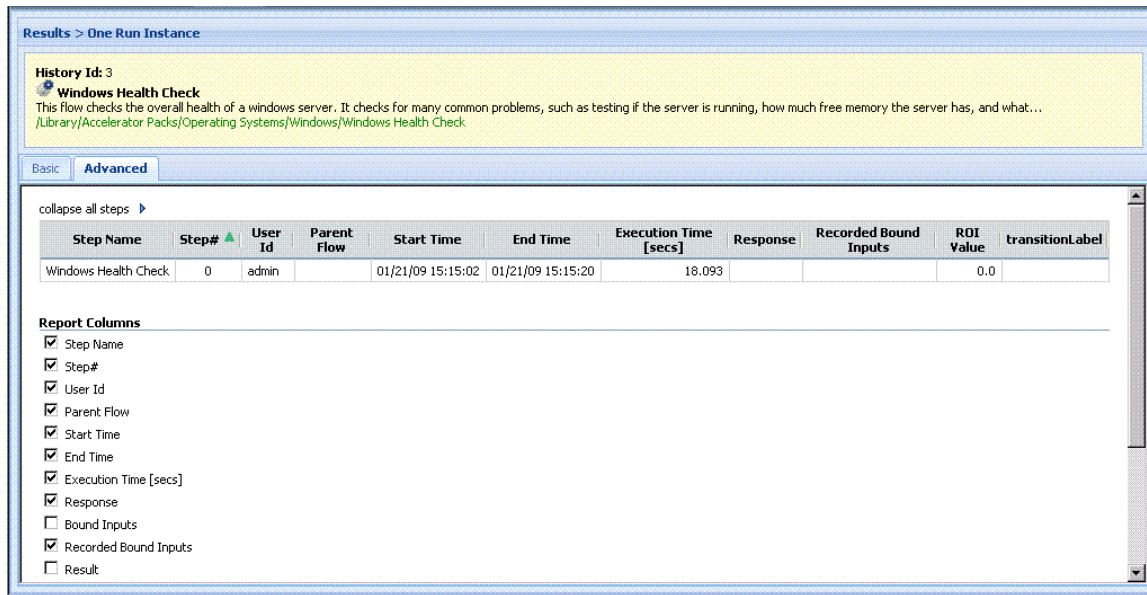


Figure 15 - Run History

The **Advanced Report** contains complete information on what has happened in each step that has completed.

You can edit which types of information the **Advanced Report** shows, or hide it altogether.

3. To select which columns appear in the **Advanced Report** section:
 - a. Scroll down below the run report.
 - b. Select the column names under **Report Columns**.
 - c. Click **Apply**.

For more information on what you can learn from run reports, see [Run histories: What happened and why](#).

4. To hide the contents of the Advanced Report, click the right-pointing arrow beside **collapse all steps**.
5. To expand the steps again, click the plus sign beside the flow and any subflows you want to expand.

Run histories: What happened and why

There are multiple reasons for examining a run in detail. For example, you might want to examine an infrastructure component such as a switch or a server that has been identified as a problem and, to do so, you might need to drill down so far as to look at the results of a single step in a single run. If your group has the RUN_REPORTS capability, you can do so on the **Reports** tab, on which you can look at run-history reports for:

- Groups of flows
- Groups of runs of one or more flows
- A single run of a single flow

On the **Reports** tab, you can search on any of the following criteria or on several of them in combination:

- The time frame within which the flow or flows were run
- Central users who ran the flows, or OO authors who created or modified the flows
- Results of or actions taken by the flows

- The flow's location in the Flow Library
- Input name/value combinations
- Configuration items or domain terms that the flows reported to the Dashboard, including the flow's category

The screenshot shows a web-based interface with a navigation bar at the top containing tabs: Dashboard, Flow Library, Current Runs, Reports (selected), Scheduler, and Administration. Below the navigation bar is a 'Filters' section with several input fields and dropdown menus. The 'Start Date' and 'End Date' fields both contain '06/18/09'. To the right, the 'User' dropdown is set to 'Any user', and the 'Results' dropdown is set to 'Any result'. The 'Location' dropdown is set to 'Any location'. Below these are empty text boxes for 'Inputs' and 'Domain Terms'. At the bottom of the filter section are 'Update Results' and 'Reset' buttons. Below the filters is a 'Results > All Flow Types' section containing a table with the following data:

Flow Name	Number of R...	Avg. Number...	Avg. Repair ...	Most Recent Run
Windows Health Check	1	38	27,111	6/18/2009 01:11:54 pm

At the bottom of the results section, there is a pagination control showing 'Page 1 of 1', a refresh icon, '50 Rows', an 'Export' dropdown, and 'Displaying 1 - 1 of 1'.

Figure 16 - Defining reports of run histories

Note: To view the **Reports** tab, your group must have the RUN_REPORTS capability.

To define a search for run histories

1. Specify the time window over which you want to see run histories. You must specify a window of time for any search to produce results.
 - For a fixed window of time:
 - a. Click the calendar in **Start Date** and select a date.
OR
Type the date in the box, in the format mm/dd/yy.
 - b. To specify a time of day for the start of the window, click TAB and then, in the drop-down box to the right of the calendar, select a time from the list. Times are shown in 12-hour format, qualified by AM or PM.
 - c. Click the calendar in **End Date** and select a date.
OR
Type the date in the box, in the format mm/dd/yy.

- d. To specify a time of day for the end of the window, click TAB and then, in the drop-down box to the right of the calendar, select a time from the list. Times are shown in 12-hour format, qualified by AM or PM.

OR

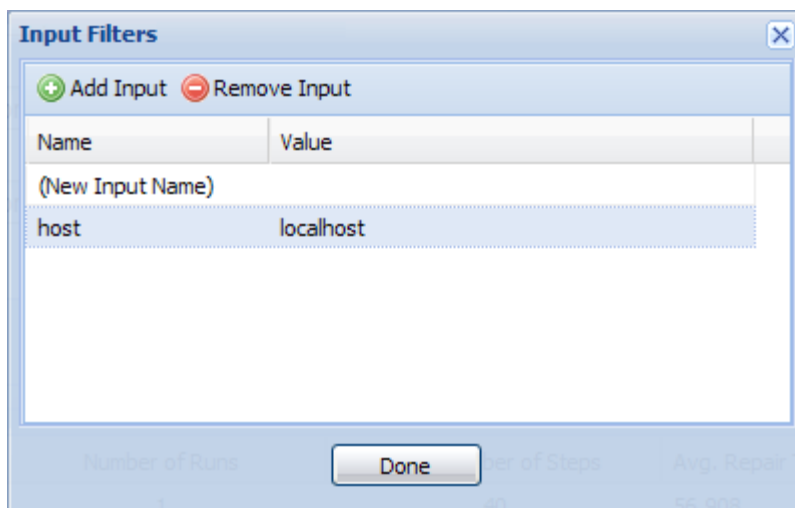
- To pick a time window relative to the present, select an item from the list under **OR choose last**.
2. To specify runs or flows related to particular users or authors, in the **User** box:
 - To specify runs or flows associated with any user, leave **Any User** in the box.
 - To specify all the runs that you have started within the time window, select **My Runs**.
To specify only runs of a certain flow or flows, you can combine this filter with a location that you specify in the **Location** box.
 - To specify all the flows whose content you have been the last person to modify, select **My flows**.
 - To specify all the flows and runs executed by one or more other users:
 1. In the **User** box, select **Run by**.
 2. In the **Run executed by users** box, type the usernames of users whose runs you want to look at, separating each two users by a comma (,).
 3. To specify one or more particular results of the flow, in the **Results** drop-down list, select the result or results that you want to include in the list of flows that appears in the report.
 4. To specify the flows in a given folder, in the **Location** drop-down box, click the Library folder that contains the flows whose runs you want to see.

The drop-down list in the **Location** box displays the Library structure of the folders that have been run.

To specify flows that have flow or step inputs of certain values, in the **Inputs** box, add any input names and values you want to use to further limit the runs included in the report:

- Click the properties icon (📄) at the right end of the **Inputs** box.

The **Input Filters** box appears:



1. To add an input, click **Add Input**.
2. Under **Name**, type the input name.
3. If you wish to specify a value for the input, then under **Value**, type the value.

OR

- To specify an input without opening the **Input Filters** box, in the **Inputs** box, type one or more inputs and their values, using the following pattern:

<name>:<value> , <name>:<value> , ...

Where

- Each <name> is the name of a different flow or step input.
 - Each <value> is the value for the input it is paired with. Note that specifying a value is optional.
 - The separator between each <name:value> pair is a comma (no space).
5. To use one or more domain terms to further define the search, then in the **Domain Terms** box, follow the same steps to specify domain term name/value pairs that you did for inputs.
 6. Click **Update Results**.

The search results appear in the **Results** list, as in the following example. In this example, the only search criterion was that the flow contain a **latencyThreshold** input.

Flow Name	Number of Runs	Avg. Number of Steps	Avg. Repair Time [secs]	Most Recent Run
Connectivity Test	1	6	52.029	2/3/2009 09:33:59 am

Figure 17 – Search Results list: a flow

The first set of results displayed in the **Results** list is always a list of the flows that have runs that fit the search criteria.

Note that the columns that are available to individual run reports are not available until you drill down to individual runs. You can, however, view a pie chart of the results of a flow in the **Results** list.

To view a pie chart of the results of a flow in the Results list

1. In the **Results** list, click the plus sign (+) next to the desired flow.
A pie chart that displays the results of the flow appears below the flow's name.

Flow Name
Connectivity Test

Results Summary

- 0% Resolved
- 0% Diagnosed
- 0% No Action
- 100% Error
- 0% Failure

The legend next to the pie chart shows the results of the flow. To help you interpret the pie chart, each result has a corresponding color. If only one of the results has a value, the pie chart shows only the color corresponding to that value.

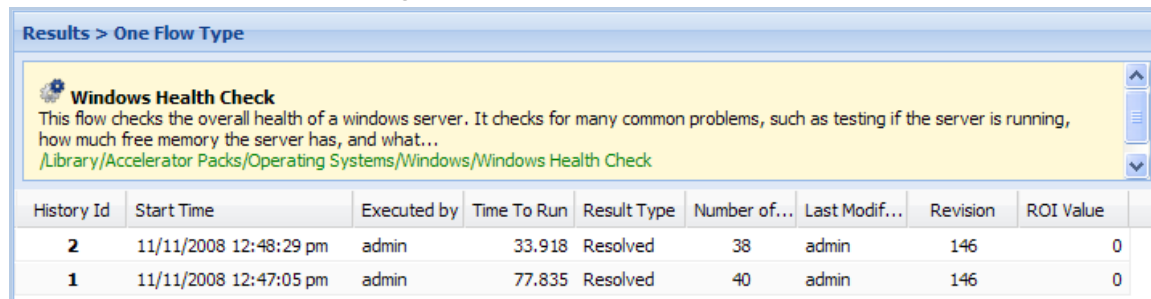
2. Click the minus sign (⊖) to close the pie chart.

Run histories in more detail

Columns that contain information on each step of a flow run provide the most detailed information on what happened in a flow run. To see more information about flow runs, including the steps in a given run, you drill down from lists of flows to individual runs, then select which columns you want to see.

To view detailed information on flow runs

1. To view the list of runs for a given flow, click the flow name.



The screenshot shows a window titled "Results > One Flow Type". Inside, there is a section for "Windows Health Check" with a description: "This flow checks the overall health of a windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what...". Below this is a table with the following data:

History Id	Start Time	Executed by	Time To Run	Result Type	Number of...	Last Modif...	Revision	ROI Value
2	11/11/2008 12:48:29 pm	admin	33.918	Resolved	38	admin	146	0
1	11/11/2008 12:47:05 pm	admin	77.835	Resolved	40	admin	146	0

Figure 18 - Runs of a flow

2. To view the steps and their information in a specific run, click the run's **History Id**.

Results > One Run Instance

History Id: 5
Windows Health Check
 This flow checks the overall health of a windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what...
[/Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check](#)

Basic | Advanced

Step	Step Name	Response	Message	Call Path
1	Ping Target System	✓	localhost successfully pinged. (0.0)	Windows Health Check > Ping Target System
2	Get System Uptime	✓	localhost has been running for SystemUptime 16587 seconds	Windows Health Check > Get System Uptime
6	Get Host Name	✓	localhost believes its name to be Name Microsoft Windows	Windows Health Check > Get Host Name
10	Get Processor Count	✓	localhost has 2 processors. They were all online . Their loads are: 0.00, 0.00, 0.00	Windows Health Check > Get Processor Count/Status
11	Free Physical Memory	✓	localhost has FreePhysicalMemory 519484 kilobytes free	Windows Health Check > Free Physical Memory
15	Get Running Processes	✓	localhost has the following processes running: Name System Idle	Windows Health Check > Get Running Process List
19	Get File System Information	✓	File system information retrieved: FreeSpace,Name,Size	Windows Health Check > Get File System Information
23	Check for Malfunctioning Devices	✓	The following devices were malfunctioning, if any: (0.7)	Windows Health Check > Check for Malfunctioning Devices
27	Get Network Interface Errors	✓	The network interface had errors (0.8)	Windows Health Check > Get Network Interface Errors

Page 1 of 4 | 10 Rows | Export | One Flow | Displaying steps 1 - 10 of 40

Figure 19 - A single run, with its steps shown

You can specify which information is shown.

To select which columns are displayed in the report of a single run

1. Put the mouse on a column heading, and then click the downward-pointing arrow at the right of the column.

Message	Call Path
localhost successfully pinged. (0.0)	
localhost has been running for SystemUptime 16587 seconds	
localhost believes its name to be Name Microsoft Windows	> Get Host Name
localhost has 2 processors. They were all online . Their loads are: 0.00, 0.00, 0.00	Windows Health Check > Get Processor Count/Status
localhost has FreePhysicalMemory 519484 kilobytes free	Windows Health Check > Free Physical Memory
localhost has the following processes running: Name System Idle	Windows Health Check

Sort Ascending
 Sort Descending
 Columns
 Step #
 Step Name
 Response
 Message
 Call Path

Figure 20 - Selecting columns to display

2. Point to **Columns**, then in the list that appears, select the columns you want to see.

Viewing step results and other advanced data

To see more information, such as the details of each step's results, you open the **Advanced View** of the run.

To view advanced data for the steps of a flow run

1. With the flow run's report open on the **Reports** tab, click the **Advanced** tab.

The **Advanced** report appears:

The screenshot shows a software interface for viewing flow run results. At the top, it says 'Results > One Run Instance'. Below that, a yellow box contains 'History Id: 5' and a gear icon next to 'Windows Health Check'. A description follows: 'This flow checks the overall health of a windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what...' and a path: '/Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check'. Below this is a tabbed interface with 'Basic' and 'Advanced' tabs, with 'Advanced' selected. A 'collapse all steps' link is visible. The main area is a table with columns: Step Name, Step#, Parent Flow, Start Time, and Response. The table contains three rows of data.

Step Name	Step# ▲	Parent Flow	Start Time	Response
<input type="checkbox"/> Windows Health Check	0		02/02/09 15:34:34	(password=(vitalProces(host=local
Ping Target System	1	Windows Health Check	02/02/09 15:35:02	success (localhost successfully pinged.) (host=local
<input type="checkbox"/> Get System Uptime	2	Windows Health Check	02/02/09 15:35:07	success (localhost has been running for SystemUpTime 16587 seconds.) (format=cs Win32_Perf

Figure 21 - Run report, advanced data, for a single flow run

You can select which data are displayed. For instance, the step results, which tend to be very long, might not be important to your purposes for this report.

2. To select which data are reported, under **Report Columns**, select the boxes for the data that you want to see, unselect those you don't, and then click **Apply**.

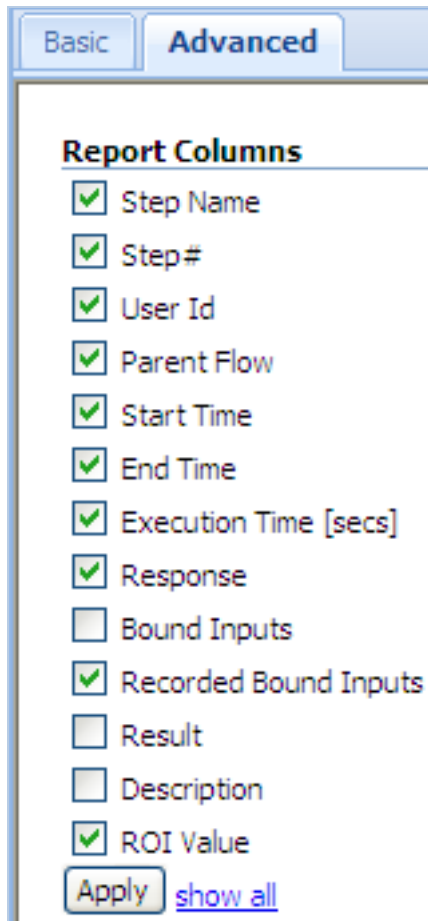


Figure 22 – Select report columns

Scheduling flows

Suppose you need to regularly check whether a number of servers are online, you can schedule a flow (say, “Connectivity Test”) to start automatically at regular intervals that you define. Each schedule that you create can specify a different server’s IP address for the flow to check. Creating schedules like this saves creating multiple flows to do the same thing, and saves you the work of starting each run individually.

There are a couple of requirements for scheduling a flow:

- The flow must be able to run automatically—that is, without requiring input from the flow user. This means that any data that the flows require must either be specific, unchanging values or be stored in flow variables, which are variables that flow authors create in Studio. When you create a schedule, you can specify input values using these flow variable names.

For example, suppose that in Connectivity Test, the flow variable **host** stores the IP address of the machine from which the flow executes the ping command. On the **Inputs** tab of the box in which you create the schedule, you would supply the IP address of the server you’re interested in. For each subsequent schedule that you create for this flow, you would specify a different IP address.

- **Necessary capabilities:** The account with which you are logged in must be a member of a group that has the SCHEDULE capability. Without the SCHEDULE or VIEW_SCHEDULES capability, you cannot see the **Scheduler** tab. If your group only has the VIEW_SCHEDULES, capability, then the flows that you see on the **Scheduler** tab are read-only for you.

Run-scheduling concurrency

You can schedule multiple runs of the same flow to run at the same time. This means that you can start multiple runs of the same flow and target them to different servers, scheduling them to all start at the same time or to start a second run of the flow before the first one ends.

Important: Suppose, however, that you schedule a flow such as a health check, to run twice against the same server, separating the two flows by a certain period of time. If one of the runs goes beyond the start time of the health check's next scheduled run, then the execution of the second run can interfere with the execution of the flow in the first run. Thus you need be aware of the possible interactions between concurrent runs of a flow. In some situations, you may need to disable run concurrency.

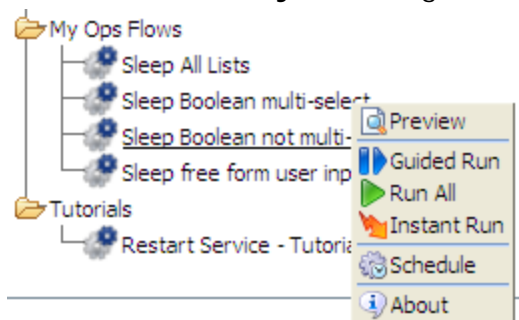
For information on disabling run concurrency see the *HP OO Administrator's Guide* (AdminGuide.pdf).

Creating a schedule for a flow

Necessary capabilities: To create a schedule, your group must have the SCHEDULE capability. Without the SCHEDULE capability, the right-click menu described in this procedure does not include the **Schedule** command, and you cannot create a flow schedule.

To create a schedule for a flow

1. On the **Flow Library** tab, navigate to and right-click the flow...



...and then, in the drop-down menu, click **Schedule**.

The **Schedule Flow** dialog box appears, with the **Schedule** tab showing.

Schedule Flow

Schedule | Inputs

Flow

Flow Name: /Library/Accelerator Packs/Operating Systems/Red Hat/Red Hat Health

Recurrence pattern

Daily
 Every : 1 minute

Weekly
 Every day

Monthly
 Every weekday

Yearly

Range of recurrence

Start date: 06/03/09
 No end date

Start time: 12:00 AM
 End after : 1 occurrences

End by : 06/03/09

Save Schedule Cancel

Figure 23 - Scheduling a flow's runs

2. On the **Schedule** tab, specify the time(s) that you want the flow to run in the boxes.
3. To specify values for the flow-level inputs, click the **Inputs** tab.

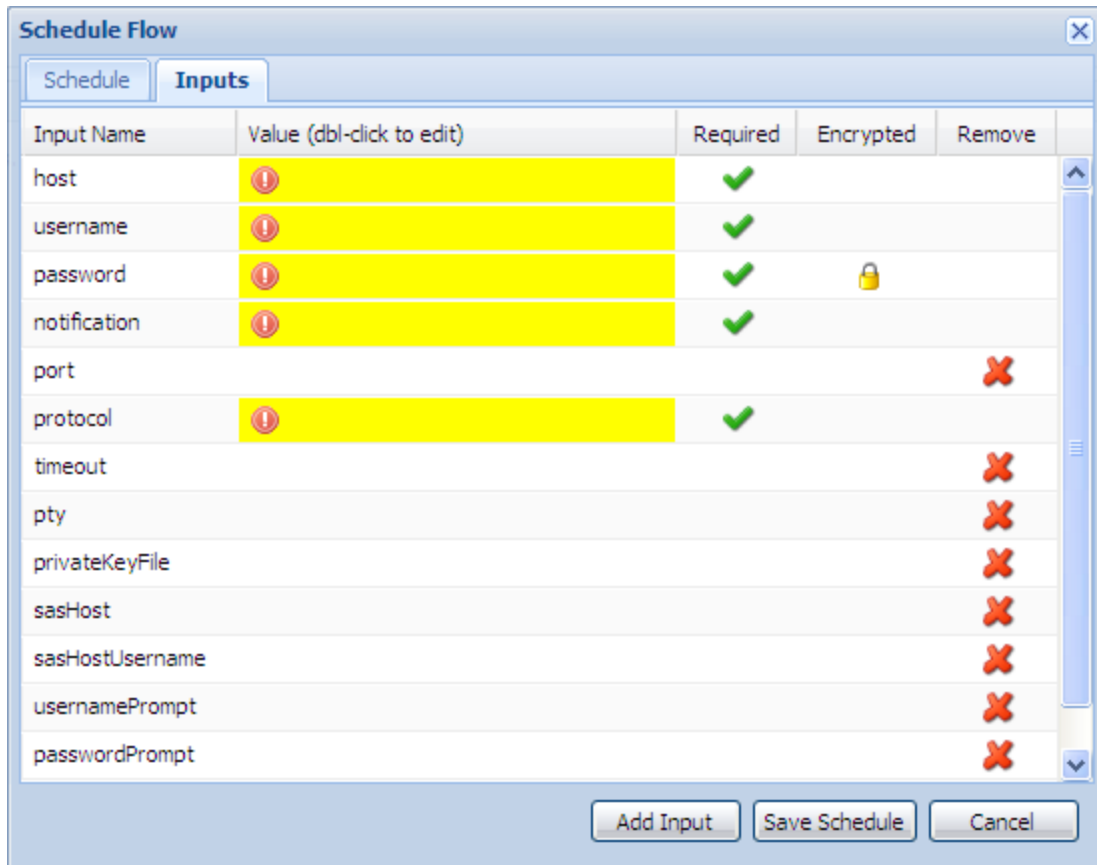


Figure 24 - Flow Scheduler Inputs tab

The flow's inputs are automatically listed in the **Input Name** boxes. If an input requires a value, the **Value** box is yellow, and there is the check mark in the **Required** column.

4. To specify a value for an input, double-click the input's **Value** box, and then do one of the following, according to whether the value is typed or chosen from a list:
 - Type the value.
OR
 - If the **List Input** dialog appears, then click the downward-pointing arrow, select the desired value from the in the drop-down list, and then click **OK**.
5. To specify multiple values for an input for a multi-instance step, type all the values in the input's **Value** box, separating them with the separator character that the flow's author specified.


For example, suppose the values for host are two IP addresses, 10.0.0.100 and 10.0.0.101 and the separator character is a comma (.). In the **Value** box, you would type:
10.0.0.100,10.0.0.101

Note that you don't include a space between the values unless the defined separator character is a sequence that includes a space.

You can remove inputs that are not required.

6. To remove an input, click the red X (✗) in the input's row under **Remove**.
7. To add an input that is not required as a flow-level input:
 - On the **Input** tab, click **Add Input**.
 - In the **Add New Input** box that appears, type a name for the input, and then click **OK**.
 - To specify a value for the input, in the table of inputs, double-click the **Value** box for the input, and then type the value (or values) for the schedule to use for the flow's runs.
8. On either tab, click **Save Schedule**.

The following appears.



Once you have created a schedule, you can edit it on the **Scheduler** tab.

Working with flow schedules

The main page of the **Scheduler** tab lists the flows that have schedules created for them.

Necessary capabilities: To view the **Scheduler** tab, your group must have either the SCHEDULE or VIEW_SCHEDULES capability. To edit a flow schedule, your group must have the SCHEDULE capability.

Scheduled Flow	Enabled	Controls	Previous run time	Next run time	Delete
Windows Health Check /Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check		Disable		Tue Nov 25 14:00:00 PST 2008	
Web Server Health /Library/Accelerator Packs/IIS/Health Check/Web Server Health		Disable		Sat Nov 29 00:00:00 PST 2008	

Figure 25 - Scheduled flows on the Scheduler tab

To see the schedules for a single flow, click the flow name. You can see the inputs that were specified for each schedule in the **Parameters** column.

Edit Schedule	Starting Date/Time	Ending Date/Time	Recurrence	Enabled	Controls	Previous run time	Next run time	Parameters	Delete
	Tue Nov 25 00:00:00 PST 2008	Wed Nov 26 00:00:00 PST 2008	every 2 hours		Disable	Tue Nov 25 14:00:00 PST 2008	Tue Nov 25 16:00:00 PST 2008	host=25.16.0.8	
	Tue Nov 25 00:00:00 PST 2008		day 1 of every month		Disable		Mon Dec 01 00:00:00 PST 2008	host=28.0.0.10	


Figure 26 - Schedules for a single flow

Note: You can edit the schedule or inputs by clicking the clock icon () in the **Edit Schedule** column.

Editing existing schedules

Note: To edit a flow schedule, your group must have the SCHEDULE capability.

To edit a flow's schedule

1. To edit a schedule, on the **Scheduler** tab, select the schedule, and then click the clock icon () in the **Edit Schedule** column.

The **Schedule Flow** dialog box appears.

2. Change the particulars of the flow schedule as desired.

For information on working in the **Schedule Flow** dialog box, see [Creating a schedule for a flow](#).

Enabling and disabling existing schedules


Note: To enable or disable a flow schedule, your group must have the SCHEDULE capability.

To enable or disable a schedule or all the schedules for a flow

1. To enable or disable all the schedules for a flow, on the **Scheduler** tab, select the flow, and then click **Enable** or **Disable** in the **Controls** column as appropriate.
2. To enable or disable a single schedule, on the **Scheduler** tab, select the schedule, and then click **Enable** or **Disable** in the **Controls** column as appropriate.

Deleting schedules

To delete a schedule or all the schedules for a flow

1. To delete all the schedules for a flow, on the **Scheduler** tab, select the flow, and then click the red ball () in the **Delete** column.
2. To delete a single schedule, on the **Scheduler** tab, select the schedule, and then click the red ball in the **Delete** column.

Configuring Scheduler settings

You can control several aspects of how the Scheduler operates by specifying settings in the **Scheduler Settings** area of the **System Configuration** subtab on the **Administration** tab. Only members of the ADMINISTRATOR group or of a group that has the MANAGE_CONF capability can see the **System Configuration** subtab.

To specify Scheduler settings

1. On the **Administration** tab, click the **System Configuration** subtab, and then scroll down to the **Scheduler Settings** area.

The area looks like this:

▼ Scheduler Configuration

Scheduler Settings

Description	Value
Log entry pattern	%-5p %d{MM/dd/yyyy HH:mm:ss} - %m%n
How many log files are retained	4
Maximum log file size. Use KB,MB,GB	10MB
The size of most recent logging sent to the UI; if the log file is bigger than this size, only this amount from the end of the file is sent. Use KB, MB or GB	64KB
If schedules of the same flow are allowed to run in parallel.	true
The account that is used to run scheduled flows. It has to be an internal account that has HEADLESS_FLOWS capability.	admin
The frequency (in ms) at which this instance "checks-in" with the other instances of the cluster. Affects the rate of detecting failed instances.	20000
If Scheduler(s) are clustered	true

Save Scheduler Changes

Figure 27 - Configurations for the Scheduler

2. Make changes in the settings as needed, taking into account the following:

Setting	What the Value changes	Notes and warnings
Log entry pattern	Enables you to configure the date, time, log message, and logging level.	Do not change the Log entry pattern unless you have a good understanding of the specifics of working with such patterns.
How many log files are retained	Limits the number of logs that are stored on the system.	Limits the amount of space needed for log data storage.
Maximum log file size	Limits the amount of space needed for log data storage.	
The size of the most recent logging sent to the UI...	Limits how much log data is presented in the Scheduler UI.	
The account that is used to run scheduled flows		By default, this is the "admin" account.
The frequency (in ms) at which this instance "checks in" with the other instances of the cluster...		Because the Scheduler is entirely part of Central, this is redundant with the clustering configuration. Take no action.

Setting	What the Value changes	Notes and warnings
If Scheduler(s) are clustered		Because the Scheduler is entirely part of Central, this is unnecessary. Always leave set to Take no action.

The Dashboard: Learning more from flows

Suppose you've had several flows running, perhaps on various schedules and using various values. The flows have been resolving alerts (or incidents, or trouble tickets), checking system and application health, and running routine maintenance on servers and applications.

Question: How can you learn the most about your infrastructure from all the work that these flows have done?

Answer: With **Dashboard** reporting charts and run-history reports on the **Reports** tab.

Note: Your group must have the RUN_REPORTS capability to view the **Dashboard** or **Reports** tab.

Collating data on Dashboard reporting charts

The **Dashboard** tab (the default starting point for Central) is a highly customizable information center for viewing the data that flows generate and analyzing the data with slices that you specify.

Note: Your group must have the RUN_REPORTS capability to view the **Dashboard** tab.

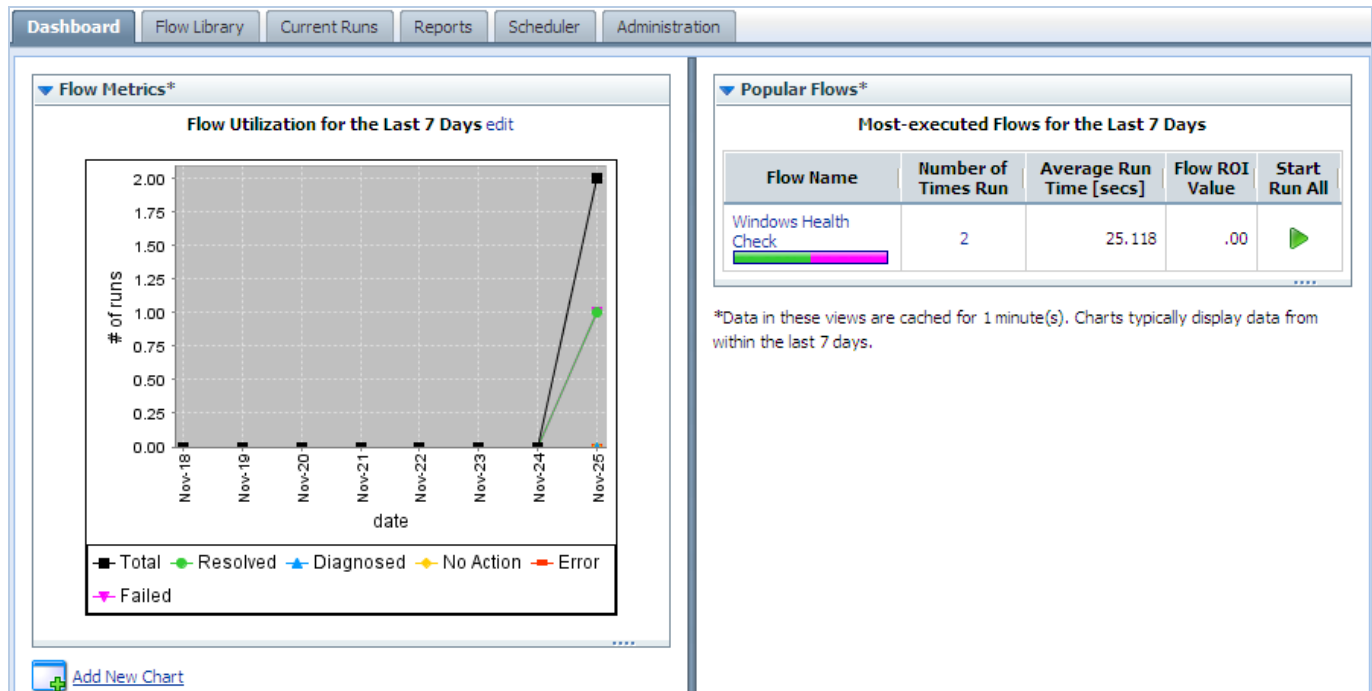


Figure 28 - The Dashboard

Charts that by default are available on the **Dashboard** tab or that you create can answer such questions as:

- Which alerts are showing up most for each application and server?

Note: OO uses the ITIL term Configuration Item (or CI) to refer to servers, applications, and other items in your operations.

- Which server or application generated the most alerts?
The **Alerts per Configuration Item** chart answers both those questions.
- What actions have been taken on each application and server?
Look at the **All CI's organized by Action** chart.
- Which flows have run the most often, and what were their outcomes?
Consider the **Outcomes per Flow** chart.
- Which flows were run to resolve errors, and how many times did the flows run?
Open the **All Alerts of Severity=Error Resolved by Flows** chart.

Most of these charts require that the flows whose data the charts report include inputs that have been configured to report their values to the domain terms that the charts use. For instance, for charts that relate alerts or actions to configuration items (CIs) to obtain enough data to work with, an input must be configured in the flow to report its value as a Configuration Item.

To display a reporting chart with current data

1. On the Dashboard, click **Add New Chart**.

The **New Chart** panel opens.

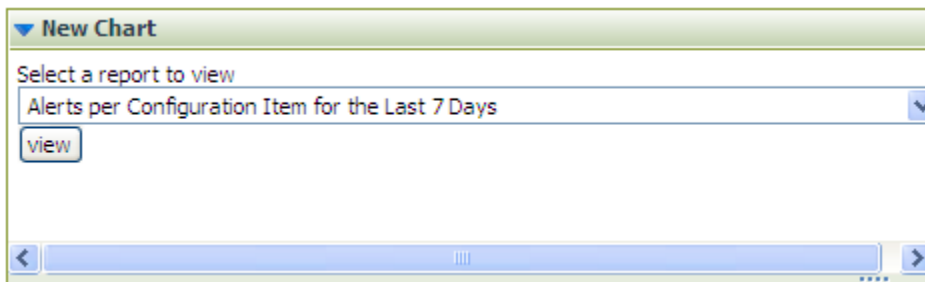


Figure 29 - New reporting chart panel

2. Select a chart from the **Select a report to view** drop-down list and click **view**.



Tip: The following are some of the domain terms that the charts record, and what they mean:

Configuration Item (aka CI)

A configuration item (CI) is any item within your infrastructure, such as a server or application. You can further categorize your CIs with CI Types and CI Minor Types. This scheme is flexible enough for you to describe the elements in your infrastructure uniquely, as the following two examples show:

A Web server:

- **CI:** the Web server's IP address
- **CI Type:** "Server"
- **CI Minor Type:** "Windows" (the Web server's operation system).

Your company home page:

- **CI:** the home page's URL
- **CI Type:** "Application"
- **CI Minor Type:** "Web Page"

Categories

The groups to which flow authors assign flows. Charting categories enables you to view performance of these classes of flows. See Studio Help for more information.

Alerts, Incidents, Problems

Alerts are monitoring messages about possible error states that have arisen amidst IT operations.

Incidents can represent trouble tickets in Incident Management or trouble-ticketing systems that you run.

Problems can represent items in any Problem Management system you operate.

Actions

What the flow did to diagnose or solve a problem or to perform maintenance, such as rebooting a server, restarting a service, changing a configuration file, re-imaging a computer, pushing new content to a Web site, or adding a new server to a cluster in order to rebalance the load.

Outcomes

Outcomes are the return states of flows: Resolved, Diagnosed, No Action, Failure

What do the bars tell you?

Let's say you're running flows that produce the following chart. This chart shows you the outcomes, whose colors align with the flow return steps whose outcomes they represent—**Diagnosed** (blue), **No Action Taken** (yellow), **Resolved** (green), and **Failure** (red).




Figure 30 - Sample Dashboard chart

The following example shows a chart that shows the actions taken per configuration item for all the flows that have run in the time specified. This is a composite that collects all the tooltips that you'll see when you move the cursor over the bar. The bar colors are generated arbitrarily when you create the chart, but are consistent within the chart.

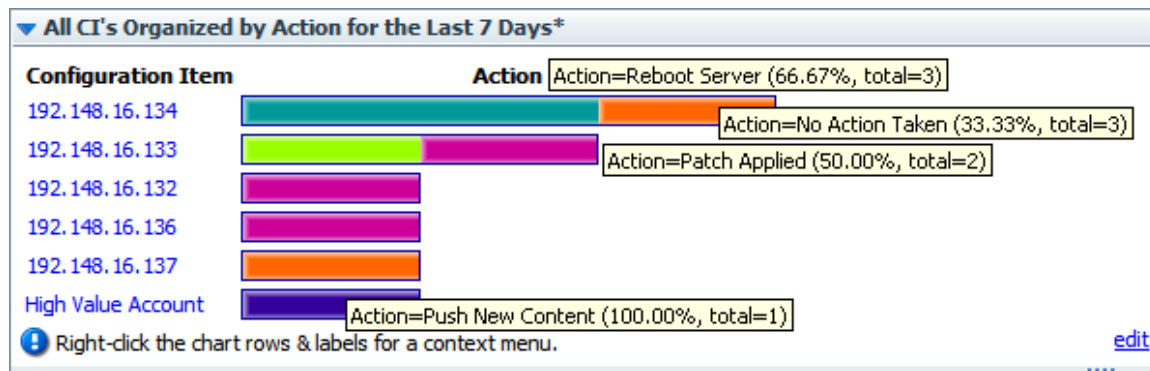


Figure 31 - Sample Dashboard chart with composite of bar labels

Each row of the chart is labeled with a configuration item's name (in this case, the server IP address or application name), but what about those bars? To learn what each bar in a row represents, float your cursor over the bar. Note that the tool tip that appears tells you:

- The action that the bar represents.
- The total number of times that that action was performed for that application or server.
- What portion of the total number of actions this particular action made up.

In this chart, we learn, among other details, that:

- Server 192.148.16.134 was rebooted by one flow and had no action taken by another.
- The Web application **High Value Account** had its content updated.

We can go further by drilling down into individual bars.

Learning more from the charts

You know what actions were performed on server 192.148.16.134, but what more can we learn about each of the actions?

To discover more, you can drill down into the chart. In the **All CI's Organized by Action** chart, let's explore the actions taken for the server 192.148.16.134.



For instance, how many alerts of what level of severity occurred that were corrected by the **Restart Service** action (the teal bar) charted here?

To learn more about data items in a chart

1. Right-click the appropriate bar segment for the data you're interested in, then click **Drill Down**.



Tip: You can also right-click the label at the left of the chart to drill down on all of the bars at once.

For instance, to show the distribution of alerts and severity levels for the Reboot Server action, right-click the bar segment representing Reboot Server, and then click **Drill Down**.

The following box appears.




Figure 32 - Creating a drill-down report

- Select a domain term for the **X** (horizontal) **Axis** and one for the **Y Axis**, then click **View**. To learn how many alerts there were of each level of severity, pick **Alert** for the **X axis** and **Severity** for the **Y axis**. This produces the following chart.

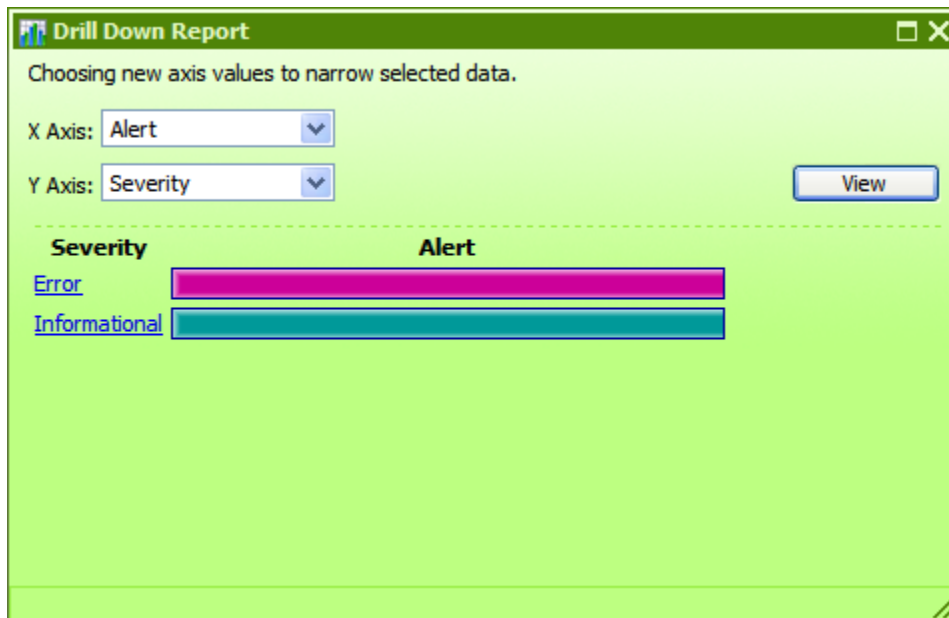


Figure 33 - The drill-down report created by choosing new axes

- To create other views of the same data, select different domain terms for the **X Axis** and **Y Axis**, and then click **View**.

As in the top-level chart, floating the cursor over a row tells you more information, such as what the alert alerted us to and how many alerts there were of this type.

You can also access the relevant run-history reports directly.

- To see a run-history report for flows that generated the data charted by the left-hand column of the chart, click the bar or the name of the row for which you want to see the run reports.

The run report is the product of a search whose terms include all the data that produced this particular bar.

In our example, if you click the **Error** row label or bar, you get a report listing all the flows that were started by alerts of severity **Error**.

To explore run-history reports, see [Run histories: What happened and why](#).

If you want information that is not charted on the charts that are available by default, you can create your own charts.

Creating and modifying charts

You can make custom charts to answer questions of your own making, such as:

- Which applications and servers had fatal errors?
- How many alerts of various severities were there?
- How many alerts were there of each kind of severity (**Informational**, **Warning**, **Error**, **Critical**, **Fatal**)?
- How many alerts of **Fatal** severity were there for each server and application?

You redefine existing charts or create new charts on the **Administration** tab, by specifying which information is charted on the horizontal (X) and vertical (Y) axes.

Note: To create or modify Dashboard charts, your group must have the `MANAGE_CONF` capability.

To create a Dashboard chart

1. On the **Administration** tab, click **Dashboard Chart Definitions**.

The page changes to show the existing chart definitions, as below.

Account Manage Users Manage Groups Dashboard Chart Definitions System Configuration

Refresh Save All Edits Create New Chart Definition Show SQL Export/Import Definitions

Alerts per Configuration Item for the Last 7 Days

*Title: Alerts per Configuration Item for the Last 7 Days Enable

*X axis: Alert X description: A monitoring alert/event ID

*Y axis: Configuration Item Y description: Machine, router, or any object with an IP address

Time Window: Last 7 Days

» Advanced Details

Save Delete

All Alerts of Severity=Error Resolved by Flows in the Last 7 Days

*Title: All Alerts of Severity=Error Resolved by Flows in the Last 7 Days Enable

*X axis: Flow X description: The name of a flow

*Y axis: Alert Y description: A monitoring alert/event ID

Time Window: Last 7 Days

» Advanced Details

Save Delete

Figure 34 - Administration tab, Dashboard chart definitions

2. Click **Create New Chart Definition**.

A new chart definition box appears.

Add New Item Here

*Title: Enable

*X axis: Action X description:

*Y axis: Action Y description:

Time Window: Last 24 Hours

»Advanced Details

Save Delete

Figure 35 - Creating a new Dashboard chart

3. Type a title for the new chart.
4. In the **X axis** drop-down list, select what you want to chart on the horizontal axis, and then type a description.
5. In the **Y axis** drop-down list, select what you want to chart on the vertical axis, and then type a description.
6. In the **Time Window** drop-down list, choose the time period you want the charting to cover – yesterday? the last week? the last month?
In the **Advanced Details** section, you can refine your charting by restricting what is charted.
7. To open the **Advanced Details** section, click **Advanced Details**.

Figure 36 - Chart definition, Advanced Details

For instance, if you want to see only certain values on the X axis, you can restrict the X axis of the chart to those values.

8. To chart only the most common occurrences of the element you're charting on the X axis or Y axis, type a number in **Top X** or **Top Y**.
For instance, to have this chart show you only the three most common types of alerts, type **3** in the **Top X** box.
9. To establish a floor value below which the X axis element is not charted, type the floor value in **X Threshold**.
For example, to leave uncharted any alert types that don't have at least five instances reported, type **5** in the **X Threshold** box.
10. To chart only elements of a certain type (as represented by a domain term value), type the domain term value in the **Restrict X axis values** box.
So suppose you want to chart only alerts of the **Loss of Connectivity** type. Assuming that the flow author has created a domain term for **Loss of Connectivity**, you could type **Loss of Connectivity** in the **Restrict X axis values** box.
Besides restricting the Y axis to the most common occurrences of the element charted there, you can further restrict what is charted in the Y axis.
11. To group any values outside those you have specified into a single color segment of the bar on the X axis, select the **Group all other values in one segment** box.
12. If you have grouped into a single X-axis segment values outside those you have specified, name the segment by typing a name in the **Name for all other values segment** box.
13. Under **Additional Constraints**, from the **Domain Term Name** drop-down, select the domain term charted on the Y axis, and then type a value in the **Domain Term Value** box.
14. In the bottom of the chart definition box, click **Save**.

Let's look at our examples:

- How many alerts are there of each kind of severity?
 - For the X axis, select Alert.
 - For the Y axis, select Severity.
- How many alerts of Fatal severity were there for each server and application?

Servers and applications are covered by the "CI Type."

- For the X axis, select Severity.
- For the Y axis, select CI Type.
- Under **Advanced Details**, in the **Restrict X Axis Values** text box, type **Fatal**.

To edit a Dashboard chart definition

1. On the **Administration** tab, click **Dashboard Chart Definitions**.
2. Scroll down to the chart you want to change.
3. In the box that defines the chart, make any desired changes as described in the above procedure [To create a Dashboard chart](#), then click **Save**.

Notes:

- Using these charts requires that the flows reported have their relevant inputs configured to report data in the domain terms that the charts need. For information on how to add this reporting capacity to inputs, see Help for Studio.
- You can add new terms that you want to appear in the **X Axis** and **Y Axis** drop-down lists. To learn more about this see the material about domain terms in Help for Studio.

Exporting and importing chart definitions

Chart definitions are persistent user reports, and are exported in an XML format. To exchange them between Central users, the users must export and import the collection of report definitions.

You can copy chart definitions from one installation of Central to another if, for instance, you want to creating consistent reporting and metrics for IT status across multiple organizations or domains. Note, however, that you do not need to export chart definitions across nodes of a Central cluster (a set of Central installations that are defined as part of a cluster).

Warning: Exporting and importing chart definitions applies to all the definitions. Any reports in the destination that have the same name as any in the XML file being imported are overwritten. There is no conflict resolution between versions of definitions that have the same name.

Note: To see the **Administration** tab's **Dashboard Chart Definitions** subtab, your group must have the `MANAGE_CONF` capability.

To export chart definitions

1. On the **Administration** tab of Central, click **Dashboard Chart Definitions**, and then **Save All Edits**.
2. Click **Export/Import Definitions**.



3. In this dialog, click **Export**.
The **File Download** dialog box appears, for the XML file that stores the chart definitions.
4. In the **File Download** dialog, click **Save**.

5. In the **Save As** dialog box that appears, navigate to the location where you want to store the XML file, and then, if you want to rename it, specify another name.
6. In the **Download complete** dialog box that appears to tell you that the export has finished, click **Close**.

OR

Open the file to examine the definitions.

To import chart definitions

1. On the **Administration** tab of Central, click **Dashboard Chart Definitions**, and then **Save All Edits**.
2. Click **Export/Import Definitions**.



3. In the **Import** text box, type a path and filename for the chart-definition XML file that you want to import.

OR

Click **Browse**, and in the **Choose File** dialog navigate to the XML file of definitions and click **Open**.

4. Back in the **Export/Import Definitions** dialog, click **Load File**.

Starting a flow from outside Central

There are a number of ways to start a flow from outside the Central application. You can:

- Use a URL created in Central to start a flow run in a Web browser either in guided-run mode (in which the flow user triggers each step manually) or in run-all mode (in which the flow runs to completion without further Central user input, apart from responses to any user prompts). To use a URL created in Central, click the flow in the **Flow Library** tab, and then copy the **Guided Run** or **Run All** URL in the **Execution Links** pane.
- Use the command-line tool RSFlowInvoke.exe, or the Java version JRSFlowInvoke.jar.
- Use the command-line tool Wget.
- Build a URL that you paste into a browser, use from a command-line tool, or use in an external system or application that can accept a command. These URLs access the REST service, allowing you to find and run flows using the Internet.
- Use the WSCentralService SOAP API to access Central features programmatically.

These techniques are described in the *HP OO SDK Guide* (SDKGuide.pdf).

Necessary capabilities: The account under which you run a flow from outside Central must be a member of a group that has the HEADLESS_FLOWS capability.

Interrupting flow runs

Keep in mind the difference between interrupting and deleting a run:

- When you interrupt a run, the progress of the run is suspended, but the run (the instance of the flow) is preserved and can be resumed.
- When you delete a run, all information about the run is deleted. You cannot restart a deleted run; you can only start a new run of the flow.

To interrupt or stop a flow run

1. During a guided or run-all run of a flow, in the upper-left panel of the flow run page, click the **Options** button (⌵).
2. In the menu that appears, click **Stop/Interrupt run**.
The run is stopped, and Central opens the **Administration** tab.

On the Central **Administration** tab, you can:

- Reassign ownership of a run
- Resume a run
- Delete a run

Reassigning ownership of a run

There are times when you might need to reassign ownership of a run, such as when the person who started a flow run does not have the required group membership for continuing past a gated transition.

Note: To reassign ownership of a run, your group must have the `MANAGE_RUNS` capability.

To reassign ownership of a flow

1. Click the **Current Runs** tab.
2. Under **Run Administration** (and above the table of runs), click **Show all users' runs**.

The **Current Runs** table changes from **My Current Runs** to **All Current Runs**, and the **User** column appears in the table.

Run Administration							
Refresh Show my runs only							
All Current Runs							
<input type="checkbox"/>	Name	Description	History ID Run ID	Started ▲	Last Modified	User	State
<input type="checkbox"/>	Windows Health Check ▼	This flow checks the overall health of a... ⓘ	3 3	5 minutes 55 seconds ago	1 minute 0 seconds ago	admin	STOP
Delete Selected						Reassign	

Figure 37 - All Current Runs table, with User column

3. In the table, under **User**, in the row for the run that you want to reassign, type the name of the OO user account that you want to own the run, then click **Reassign**.

Resuming a run

To resume a run that you did not start and that has not been re-assigned or handed off to you, you must be a member of the ADMINISTRATOR group or of a group that has the MANAGE_RUNS capability. You resume runs on the **Current Runs** tab.

Warning: The only runs that you should resume are those that have been interrupted or handed off to you. Clicking **Resume** for a run that is currently running transfers ownership of the run from the user who started it to you. This causes an error for the user whose active run you resumed. In addition, other problems may result. Active runs include those whose state is **RUNNING**, **IDLE**, **WAITING_USER_INPUT**, or **NOT_STARTED**. For information on reassigning ownership of a run, see [Reassigning ownership of a run](#).

Notes:

- History ID and Run ID can be but are not necessarily the same for a particular run. When you click **History ID|Run ID** to sort the table of **Current Runs**, the table sorts on History ID.
- If you resume a run that is already running, it will go into a PAUSED state, and you must click either **Next Step** or **Run All** in the flow graph.

To resume a flow run

1. Click the **Current Runs** tab.




Figure 38 – Administering runs

2. Under **Name**, click the downward-pointing arrow by the name of the flow whose run you want to resume.

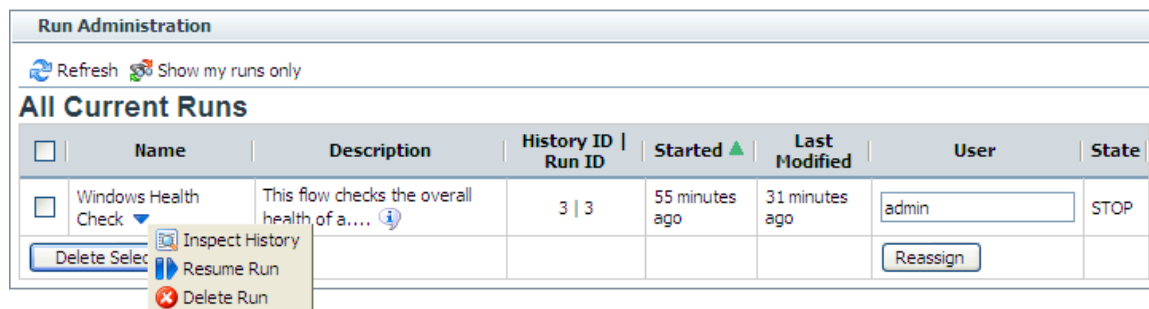


Figure 39 - Context menu for a flow run

3. On the context menu that appears, click **Resume Run**.
The run is loaded in OO.
4. Click either **Next Step** or **Run All**, according to which mode you want to complete the run in.

Deleting a run

Note: If you delete a run from the **Current Runs** tab, the History ID and the Run ID can be but are not necessarily the same for a particular run. History ID is universally unique across the history of

the system. When you click **History ID|Run ID** to sort the table of **Current Runs**, the table sorts on History ID.

To delete a run

1. Click the **Current Runs** tab.




Figure 40 – The current runs

2. Under **Name**, click the downward-pointing arrow by the name of the flow whose run you want to delete.




Figure 41 - Context menu for a flow run

3. On the context menu that appears, click **Delete Run**.

To delete a group of runs

1. Select the check box beside the name of each flow run that you want to delete.

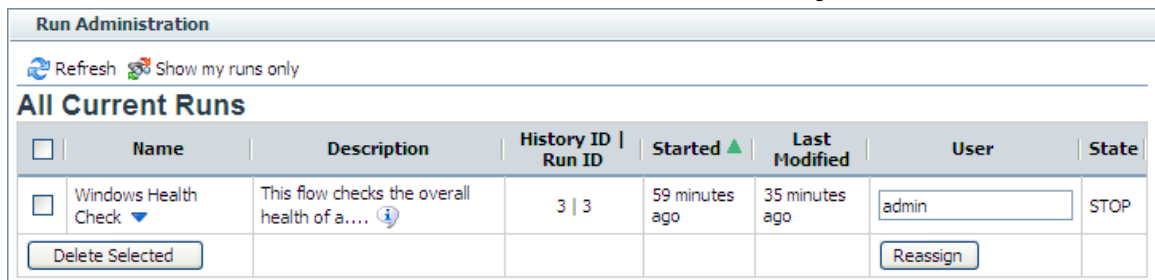


Figure 42 - Run Administration

Note: If there is more than one current run for a flow, you may need to identify the run by its History ID.

OR

To select all the current runs, select the check box in the title row of the table of runs.

2. Click the **Delete Selected** button.

Creating a link to a flow run

You can also enable another Central user to resume a flow run that you have paused (or resume the run yourself) by capturing the run’s URL mid-run, then placing the URL another Web page or in an

IM or email message. The URL that you capture is not the same as the address that is in your Web browser address box during the run.

You can create a link to either a guided run or a run-all run of the flow, but not an instant run.

To create a link to a flow run

1. When previewing a flow or during a run (except an instant run), above the run-type buttons in the upper-left panel of the flow run page, click the **Options** button (🔍).

2. In the menu that appears, click **Link to this run**.

The **Link to run** dialog box appears.

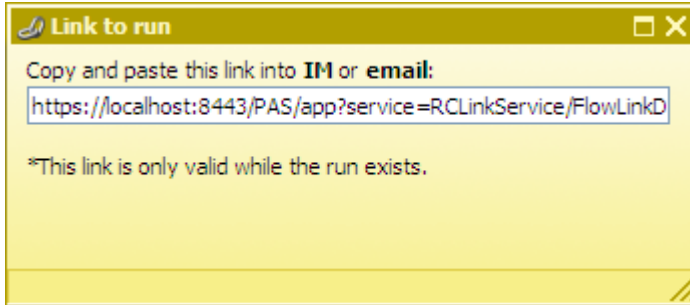


Figure 43 - Creating a link to the run

3. In the **Copy and paste this link...** text box, copy the entire URL, and then close the **Link to run** dialog box.

4. In the external source in which you are creating access to the run, paste the shortcut that you copied.

Creating a link to a flow

You can make a link to the flow available to other users by capturing its URL, then placing the URL on another Web page or in an IM or email message. The URL that you capture is not the same as the address in your Web browser address box.

To create a link to a flow

1. On the **Flow Library** tab, double-click the flow to open its preview.




Figure 44 - Preview, left pane

- At the bottom of the pane, click the **Execution Links** bar. In the **Execution Links** section that appears, links that you can paste into a Web browser's URL address line to start the flow appear, labeled for the kind of run you can carry out with them.

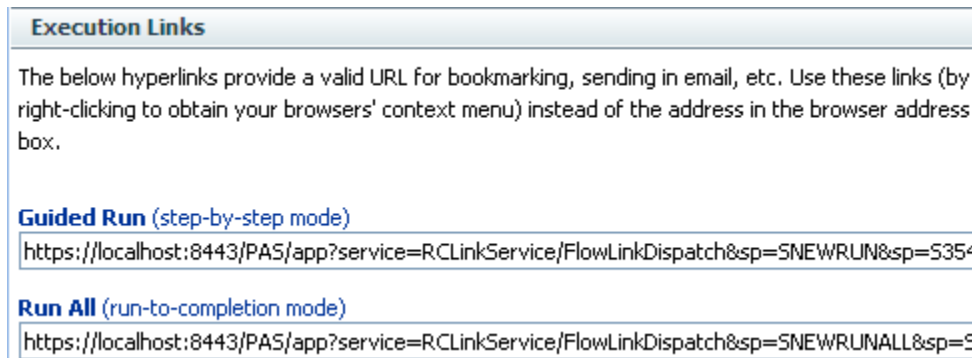


Figure 45 - Links that will run the flow

- Copy the appropriate URL, then paste it into an email message or the URL address box of a Web browser instance.

Handing off flow runs

You might need to hand off a flow in which:

- A step requires information that someone else has.
- A transition is gated (requires membership in a group that your account is not a member of).

Note: The person who resumes the run must have sufficient permissions to execute the flow.

You can hand off a guided run or a run-all run of the flow. An instant run that has a gated or hand-off transition can also be handed off.

To hand off a flow run

1. During a guided or run-all run of a flow, above the run-type buttons in the upper-left panel of the flow run page, click the **Options** button (🔍).
2. In the menu that appears, click **Hand off**.

The run is paused, and the state of the run is changed to **Handed Off**.

A new e-mail message appears, with the URL of the flow included in the body of the message.

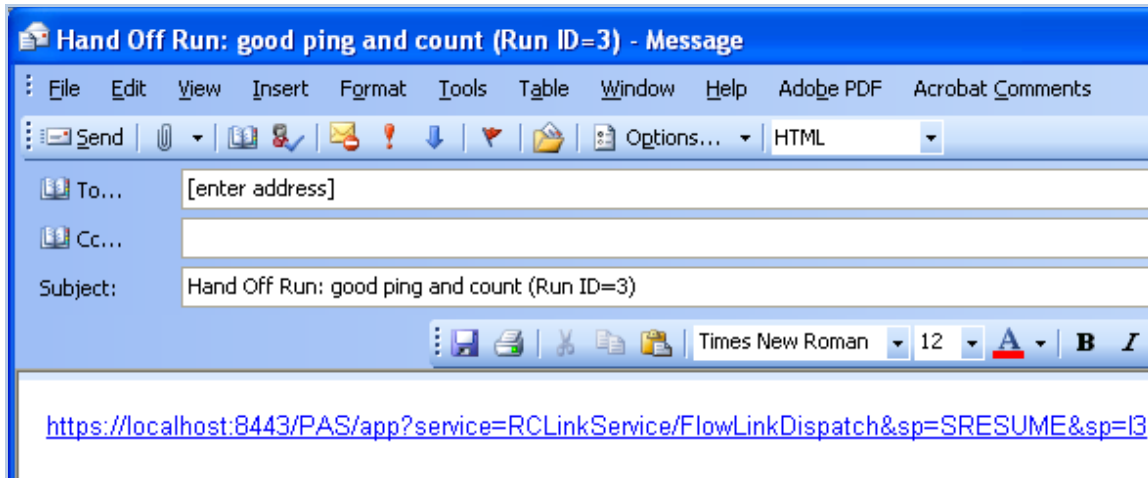


Figure 46 - Flow run handoff e-mail message

3. Address the message to the person to whom you're handing off the flow, and send the message.

To resume a flow run that has been handed off to you

1. Open the e-mail message that contains the URL of the flow and click the URL.
A new browser instance of Central opens, with the flow run ready to resume. You can get it going again as you would any other flow run that is paused.
2. To continue the flow run, click either the **Next Step** or the **Run All** button.

Auditing and managing flows

Audit information on flows (individually or groups) and their runs can be particularly important in system diagnostics at several levels, for Central users, OO administrators, and IT managers. For information on auditing and viewing reports on flow runs, see [Run histories: What happened and why](#).

While all users of Central can interrupt, cancel, resume, or restart flow runs that they own, in order to view all users' current flow runs and manage (pause, resume, or delete) them, a user must be a member of a group that has the MANAGE_RUNS capability.

Users, groups, and access control

The whole point of working with user accounts is to enable the right people to run flows in Central and (for authors) to create flows in Studio. Toward that end, you work with users, groups, capabilities, and permissions. In order, you do the following:

1. Enable OO to use the kind of authentication that your system employs.
2. Add users to OO and add them to groups or map their external groups to OO groups.
3. Grant capabilities to OO users and groups.

Individual flow authors grant access permissions to their flows. For information on assigning users or groups access permissions for various OO objects, see Help for Studio.

Because capabilities are a key concept for this way of controlling who can do what, we'll consider [Capabilities and access permissions](#) before explaining how to:

- Configure OO for working with authentication by Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or Kerberos providers, and add OO users, in [Allowing external users into the Central system](#).
- Manage users, in [Managing users](#).
- Manage groups, in [Managing groups](#).
- Setting logging levels and other system settings

For more information, see the *OO Administrator's Guide* (AdminGuide.pdf).

Capabilities and access permissions

Working with flows, schedules, users, and other OO objects requires a combination of capabilities (the ability to perform an action) and the necessary access permissions for each flow and the objects associated with it.

- A capability is the right to perform an action in OO, such as the `MANAGE_USERS` and `MANAGE_GROUPS` capabilities. A user with the `MANAGE_GROUPS` capability can assign groups the capabilities that they need. For more information, see [Capabilities](#).
- Permissions are access rights to individual objects: folders, flows, operations, or system accounts. The four permissions are Read, Write, Execute, and Link to, which flow authors grant to users or groups for individual objects. So:
 - To view and find a flow in Central, users must have Read permissions for the flow.
 - To run a flow in Central, users must have Read and Execute permissions for the flow and its subflows, and Read and Write permissions for any system accounts used by the flow.
 - In Studio, authors must have Read, Write, Link to, and/or Execute permissions for objects that they use to author flows. For instance:
 - To debug a flow, an author must have the Execute permission for that flow.
 - A flow author must have the Link to permission for any flow or operation from which he or she creates a step in a flow.
 - To change a system account, an author must have the Read and Write permissions for the system account.

For more information, see [Permissions](#).

Capabilities

Following are the capabilities that can be assigned in OO.

Capability	Description
<code>MANAGE_USERS</code>	Create, delete, and modify internal users and map external user groups (groups that exist outside of OO) to internal OO groups. Only holders of this capability can create internal OO users.
<code>MANAGE_GROUPS</code>	Create, delete, and modify groups.
<code>AUTHOR</code>	Start Studio.

Capability	Description
SCHEDULE	Schedule flows.
MANAGE_RUNS	View, start, stop, delete, and reassign runs other than the user's.
RUN_REPORTS	View reports and Dashboard metrics/charts.
MANAGE_CONF	Manage configuration properties and dashboards.
VIEW_SCHEDULES	View flow schedules.
HEADLESS_FLOWS	Start flows from outside Central.

Permissions

The following two tables describe the permissions and which of them are needed for objects in Studio.

Permissions for OO objects

Permission	Description
Read (R)	Can view the object in Studio or Central.
Write (W)	Can change the object.
Execute (X)	Can start a run of the flow. This is not a recursive requirement. That is, for a Central user to run a flow or for an author to debug a flow, he or she does not have to have Execute permission for all the objects, such as operations and configurable items, associated with the flow. The user does, however, need Read and Write permissions for the objects associated with the flow.
Link to (L)	Can use the flow or operation to create a flow step.

OO objects and the permissions needed to work with them

Object	Action	Necessary permission(s)
Folder		
	View contents	Read
	Add to contents	Read, Write (also needed for all children of the folder)
	Move	Read, Write
	Rename	Read, Write
Flow or operation		

Object	Action	Necessary permission(s)
	View/open	Read
	Modify	Read, Write
	Rename	Read, Write
	Execute/Run	Read, Execute
	Use as a step or subflow	Read, Link to
System account		
	View account name	Read
	Change account password	Read, Write
	Rename account	Read, Write
	Use in flow or operation	Read, Link to
	Use at runtime	Read, Execute

For more information on the groups, capabilities, and permissions model of OO security, see the *OO Administrator's Guide* (AdminGuide.pdf).

When you first deployed OO, you mapped users to the OO groups. Depending on how you accomplish the mapping, when you deploy the OO clients to an additional user, you add the user to a group either by adding the user to the appropriate group or role in your organization's authorization system or by individually mapping the user to an OO group. For information on mapping users to OO group, see the installation and deployment guide, *Installing or Upgrading HP Operations Orchestration* (InstallGuide.pdf).

Allowing external users into the Central system

Besides creating users within OO (called *internal* users), you can control who can use OO by mapping external users or groups to OO groups. This is very useful for creating large numbers of user accounts. You authenticate external user accounts for OO by enabling an external authentication provider—Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or Kerberos—then map the external groups to OO groups. (For information on enabling authentication providers, see [Using external authentication for Central users](#).)

However, it's not likely that you will want all the authenticated users in a large organization to use Central or Studio. So in addition to requiring authentication by an external authentication provider, OO requires that external user accounts also be mapped to an OO group because either:

- The external user's account is a member of an external group that is mapped to an OO group
- Users authenticated by the external authentication provider are automatically mapped to a particular Central group.

When an external user account is assigned to any Central group, it is always also assigned to the EVERYBODY group. This provides everyone who is allowed to use Central with the common baseline of capabilities.

Necessary capabilities: To map external users or groups to OO groups, you (that is, the account that you are logged in with) must be a member of a group that has the `MANAGE_GROUPS` capability. To configure external authentication for OO, your group needs to have the `MANAGE_CONF` capability.

Using external authentication for Central users

To authenticate with any or all of Active Directory (AD), LDAP, or Kerberos authentication providers, you use the **Administration** tab's **System Configuration** subtab. The **System Configuration** subtab contains a section for each of the three kinds of authentication providers. The following might look a bit formidable if you are not an AD, LDAP, or Kerberos administrator, but in the procedure following this illustration, we'll work through configuring the settings that are relevant for the type (or types) of authentication that your organization uses. You may need to consult with the IT administrator who configured your authentication and/or directory.

Necessary capabilities: To configure Central to use external authentication, you must be a member of a group that has the MANAGE_CONF capability.

To modify the configuration values for authentication providers according to your organization's needs, use one or more of the following sections:

- [AD authentication settings](#)
- [LDAP authentication settings](#)
- [Kerberos authentication settings](#)

For these sections on configuring OO using external authentication, suppose the following users are members of these groups:

User	Member of this external group
Tom Gage, a service desk technician	Service Desk
Mary Grey, a network specialist	Network Specialist
Ed Stuart, a system administrator	Manager

Suppose also that the name of their domain is "mirage," and suppose the following about the domain server:

- Its IP address is 192.111.5.102
- Its fully qualified name is mirage.ad

AD authentication settings

The following procedure for authenticating with AD refers to this section of the **System Configuration** subtab of the **Administration** tab.

Necessary capabilities: To configure Central to use external authentication, you must be a member of a group that has the MANAGE_CONF capability.

AD Authentication Settings AD Enabled

Description	Value
The default group an AD authenticated user gets when there is no group matching.	EVERYBODY
LDAP filters that try to match the user groups. These filters are applied to the discovered groups and if they match, the user is considered part of that group. The list separator is a ";".	member=cn={1},CN=Users,DC=MyCompany,DC=ad
List of LDAP contexts containing user groups. For example, if you have a "Groups" object containing groups of users, then the expression: ou=Groups,dc=MyCompany,dc=com might be used. The list separator is a ";".	CN=Users,DC=MyCompany,DC=ad
Attribute of any group (returned from the group search), to use as group name.	name
Active Directory URL.	LDAP://MyCompany.ad
The user domain (the authentication used is the NT style "domain\username").	MyCompany\{0}
List of LDAP contexts containing users. For example, if you have a "Users" object containing users, then the expression: ou=Users,dc=MyCompany,dc=com might be used. The list separator is a ";".	CN=Users,DC=MyCompany,DC=ad
<input type="button" value="Save AD Settings"/> <input type="button" value="Test"/> <input type="button" value="Refresh"/>	

Figure 47 - AD Authentication Settings

To authenticate with Active Directory

1. Select the **AD Enabled** check box.
2. In the **Value** box for the **default group**, specify which OO group or individual is mapped to when a mapping is not specified.
For instance, to assign unmapped groups or individuals only the capability to run flows, you would type **LEVEL_ONE**. To learn about the default OO groups, see [Managing groups](#).
3. In the **Value** box for **LDAP filters that try to match the user groups**, type a filter to find the groups for users.
For instance, suppose the following:
 - Tom Gage's "memberOf" entry has a value of "CN=Service Desk,CN=Users,DC=mirage,DC=com"
 - Mary Grey's "memberOf" entry has a value of "CN=Network Specialist,CN=Users,DC=mirage,DC=com"
 - Ed Stuart's "memberOf" entry has a value of "CN=Manager, OU=Staff, DC=mirage, DC=com"
Thus Tom's and Mary's LDAP entries are defined under the context CN=Users, while Ed's entry is defined under OU=Staff.
In the **Value** box, you might specify the following filter:
`member=CN={1},CN=Users,DC=mirage,DC=com; member=CN={1},OU=Staff,DC=mirage,DC=com`
Be sure to:
 - Type each instance of **member=CN={ 1 }**, exactly as it appears above.
 - Type the semicolon (;) separator between the filters that you type.
4. In the **Value** box for the **List of LDAP contexts containing user groups**, type the contexts in which LDAP should search for your existing, external groups.
You can provide multiple contexts, using commas to separate the relative distinguished names (RDNs) within a context and the semicolon (;) to separate contexts.

For example, suppose that:

- Tom Gage's "Service Desk" group and Mary Grey's "Network Specialist" group are defined under OU=Groups (OU=Groups,DC=mirage,DC=com).
- Ed Stuart's "Manager" group is defined under OU=Staff (OU=Staff,DC=mirage,DC=com)

The following setting makes these groups visible to OO:

```
OU=Groups,DC=mirage,DC=com;OU=Staff,DC=mirage,DC=com
```

Next we need to create search filters, to tell OO how to find groups (in roleContextsList) that point to users.

5. Leave **name** in the **Value** box for the **Attribute** setting.
6. In the **Value** box for the **Active Directory URL** setting, supply the URL or IP address of the AD, with the following syntax:

```
LDAP://<AD_server>[:<port>]
```

where:

- <AD_server> is the IP address or fully qualified name of the AD server.
- <port> is the port number that the AD server uses, if the AD server is configured to use a non-standard port (that is, other than 389). If the AD server uses port 389, you can omit :<port> from the setting.

For instance, if the AD server is mirage.ad, its IP address is 192.168.5.5, and it uses port 200, the setting would be:

```
LDAP://mirage.ad:200
```

or

```
LDAP://192.168.5.5:200
```

Important: Machines ordinarily communicate with Active Directory using Lightweight Directory Access Protocol (LDAP, a clear-text protocol). To encrypt communications, you can set OO to communicate with Active Directory over Secure Sockets Layer (SSL). The LDAPS protocol is the LDAP protocol encrypted with SSL. If you want to encrypt Active Directory communications in your organization, see the *OO Administrator's Guide* (AdminGuide.pdf) for information on configuring your system to use the LDAPS protocol.

If LDAP is configured over SSL, the protocol portion of the AD URL should be [LDAPS](#), so the setting would be:

```
LDAPS://mirage.ad:200
```

or

```
LDAPS://192.168.5.5:200
```

7. In the **Value** box for the **user domain** setting, type the domain where the users reside.
<domain>\{0}
Be sure to type **{0}** exactly as it appears above.
8. To specify the contexts in which OO should look for users, type the contexts in the **Value** box for the **List of LDAP contexts containing users** setting.

In our example, Tom's, Mary's, and Ed's entries are defined under the same contexts as the groups that they belong to are. Thus we would provide the following:

```
OU=Users,DC=mirage,DC=com;OU=Staff,DC=mirage,DC=com
```

Using this procedure's example, the AD Authentication area should look like the following:

AD Authentication Settings <input checked="" type="checkbox"/> AD Enabled	
Description	Value
The default group an AD authenticated user gets when there is no group matching.	LEVEL_ONE
LDAP filters that try to match the user groups. These filters are applied to the discovered groups and if they match, the user is considered part of that group. The list separator is a ";".	member=cn={1},CN=Users,DC=mirage,DC=ad
List of LDAP contexts containing user groups. For example, if you have a "Groups" object containing groups of users, then the expression: ou=Groups,dc=MyCompany,dc=com might be used. The list separator is a ";".	CN=Users,DC=mirage,DC=ad
Attribute of any group (returned from the group search), to use as group name.	name
Active Directory URL.	LDAP://mirage.com
The user domain (the authentication used is the NT style "domain\username").	mirage\{0}
List of LDAP contexts containing users. For example, if you have a "Users" object containing users, then the expression: ou=Users,dc=MyCompany,dc=com might be used. The list separator is a ";".	CN=Users,DC=mirage,DC=ad
<input type="button" value="Save AD Settings"/> <input type="button" value="Test"/> <input type="button" value="Refresh"/>	

Figure 48 - AD authentication enabled

9. To save your settings, click **Save AD Settings**.
10. To test the current AD authentication settings from within OO, click the **Test** button. The **Testing AD Settings** dialog box appears.

Testing AD Settings
✕

Enter a sample user to test authentication with your new settings.

User Name*

Password*

* user/password credentials supplied for testing are not saved

Results will appear here.

Figure 49 - Testing AD authentication settings

11. Type the user name and password of an external account that is authenticated by AD, and then click **Test**.
12. If the test fails, modify the settings and test again.
13. When the AD settings test successfully, click the **Save AD Settings** button.
14. When you see a message to restart the Central service (RSCentral on a Windows system or Central.sh on a Linux system), do so.

If your IT organization also authenticates with LDAP and/or Kerberos, complete the procedures in [LDAP authentication settings](#) or [Kerberos authentication settings](#). Finally, to map external Active Directory groups to OO groups, see [Mapping external groups to OO groups](#).

LDAP authentication settings

The following procedure for authenticating with LDAP refers to this section of the **System Configuration** subtab of the **Administration** tab.

Necessary capabilities: To configure Central to use external authentication, you must be a member of a group that has the `MANAGE_CONF` capability.

LDAP Authentication Settings LDAP Enabled

Description	Value
The default group an LDAP authenticated user gets when there is no group matching.	EVERYBODY
LDAP search filter that tries to match the user groups (see RFC 2254 for LDAP search filter syntax). This filter is applied to the discovered groups and if it matches, the user is considered part of that group.	((!(member=cn={1},ou=Users,dc=atlantis,dc=com)(member=cn={1},ou=Machinists,dc=atlantis,dc=com))
List of LDAP contexts containing user groups. For example, if you have a "Groups" object containing groups of users, then the expression: ou=Groups,dc=mycompany,dc=com might be used. The list separator is a ";".	OU=Groups,DC=atlantis,DC=com;OU=Machines,DC=atlantis,DC=com
Attribute of any group (returned from the group search), to use as group name.	name
LDAP URL.	LDAP://192.168.88.128
List of LDAP contexts containing users. {0} denotes the location where the username should be inserted to create a DistinguishedName. The list separator is a ";".	CN={0},OU=Users,DC=atlantis,DC=com;CN={0},OU=Machinists,DC=atlantis,DC=com
List of user context attribute names which can be used as groups. The list separator is a ";".	

Figure 50 - Settings for configuring OO with LDAP authentication

To authenticate with LDAP

1. Select the **LDAP Enabled** check box.
2. In the **Value** box for the **default group**, specify which OO group or individual is mapped to when a mapping is not specified.
 For instance, to assign unmapped groups or individuals only the capability to run flows, you would type **LEVEL_ONE**. To learn about the default OO groups, see [Managing groups](#).
3. In the **Value** box for **LDAP search filter that tries to match the user groups**, type a filter to find the groups for users.
 For instance, suppose the following:
 - Tom Gage's "memberOf" entry has a value of "CN=Service Desk,CN=Users,DC=mirage,DC=com"
 - Mary Grey's "memberOf" entry has a value of "CN=Network Specialist,CN=Users,DC=mirage,DC=com"
 - Ed Stuart's "memberOf" entry has a value of "CN=Manager, OU=Staff, DC=mirage, DC=com"
 Thus Tom's and Mary's LDAP entries are defined under the context CN=Users, while Ed's entry is defined under OU=Staff.

In the **Value** box, you might specify the following filter:

```
( | (member=CN={1} , CN=Users , DC=mirage , DC=com) (member=CN={1} , OU=Staff , DC=mirage , DC=com)
```

This filter finds groups in either:

```
member=CN={1},CN=Users,DC=mirage,DC=com
```

or

```
member=CN={1},OU=Staff,DC=mirage,DC=com.
```

Be sure to:

- Type each instance of **member=CN={1}**, exactly as it appears above.
 - If you type more than one filter, separate the filters with a semicolon (;).
4. In the **Value** box for the **List of LDAP contexts containing user groups**, type the contexts in which LDAP should search for your existing external groups.

You can provide multiple contexts, using commas to separate the relative distinguished names (RDNs) within a context and the semicolon (;) to separate contexts.

For example, suppose that:

- Tom Gage's "Service Desk" group and Mary Grey's "Network Specialist" group are defined under OU=Groups (OU=Groups,DC=mirage,DC=com).
- Ed Stuart's "Manager" group is defined under OU=Staff (OU=Staff,DC=mirage,DC=com)

The following setting makes these groups visible to OO:

```
OU=Groups,DC=mirage,DC=com;OU=Staff,DC=mirage,DC=com
```

Next we need to create search filters, to tell OO how to find groups (in roleContextsList) that point to users.

5. Leave **name** in the **Value** box for the **Attribute** setting.
6. In the **Value** box for the **LDAP URL** setting, supply the URL or IP address of the top level of the LDAP server, with the following syntax:

```
LDAP://<LDAP_server>[:<port>]
```

where:

- `<LDAP_server>` is the IP address or fully qualified name of the LDAP server.
- `<port>` is the port number that the LDAP server uses, if the LDAP server is configured to use a non-standard port (that is, other than 389). If the LDAP server uses port 389, you can omit `:<port>` from the setting.

For instance, if the LDAP server is mirage.ad, its IP address is 192.168.5.5, and it uses port 200, the setting would be:

```
LDAP://mirage.ad:200
```

or

```
LDAP://192.168.5.5:200
```

Important: To encrypt communications, you can set OO to communicate with LDAP over SSL by specifying the LDAPS protocol. For information on configuring your system to use the LDAPS protocol, see the *OO Administrator's Guide* (AdminGuide.pdf).

If LDAP is configured over SSL, the protocol portion of the AD URL should be **LDAPS**, so the setting would be:

```
LDAPS://mirage.ad:200
```

or

```
LDAPS://192.168.5.5:200
```

7. To specify the contexts in which OO should look for users, type the contexts in the **Value** box for the **List of LDAP contexts containing users** setting.

In our example, Tom's, Mary's, and Ed's entries are defined under the same contexts as are the groups that they belong to. Thus we would provide the following:

```
OU=Groups,DC=mirage,DC=com;OU=Staff,DC=mirage,DC=com
```

8. To specify one or more LDAP user context attributes for use as OO group names: In the **Value** box for the **List of user context attribute names which can be used as groups**, list the LDAP user context attributes, separating each attribute from the next with a semicolon (;).
9. To save your settings, click **Save LDAP Settings**.
10. To test the current LDAP authentication settings from within OO, click the **Test** button. The **Testing LDAP Settings** dialog box appears.

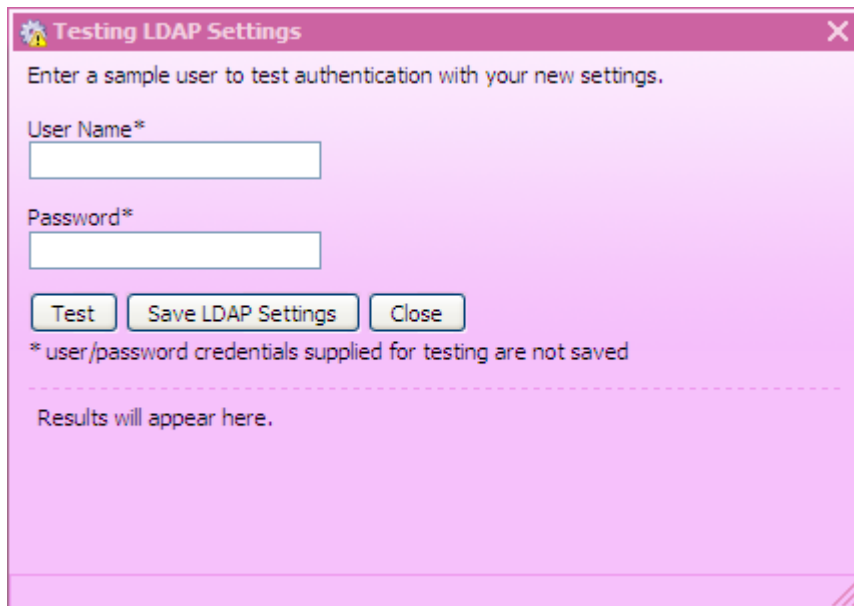


Figure 51 - Testing LDAP authentication settings

11. Type the user name and password of an external account that is authenticated by LDAP, and then click **Test**.
12. If the test fails, modify the settings and test again.
13. When the LDAP settings test successfully, click the **Save LDAP Settings** button.
14. When you see a message to restart the Central service (RSCentral on a Windows system or Central.sh on a Linux system), do so.

If your IT organization also authenticates with Active Directory and/or Kerberos, complete the procedures in [AD authentication settings](#) or [Kerberos authentication settings](#). Finally, to map external LDAP groups to OO groups, see [Mapping external groups to OO groups](#).

Kerberos authentication settings

Kerberos authenticates only individual users, so when you use Kerberos authentication, you cannot map external groups to OO groups. After you have configured the Kerberos authentication settings as necessary, you will use the **Manage Users** subtab to assign authenticated users to OO groups.

The following procedure for authenticating with Kerberos refers to this section of the **System Configuration** subtab of the **Administration** tab.

Necessary capabilities: To configure Central to use external authentication, you must be a member of a group that has the MANAGE_CONF capability.

Kerberos Authentication Settings <input type="checkbox"/> Kerberos Enabled	
Description	Value
Kerberos5 configuration file. It should be relative to product install location (e.g /Central/conf/krb5.conf). If realm and kdc host are provided, they override the default realm and KDC values from the conf file.	<input type="text"/>
KDC host (format is "address:port" or "address" when using the default port).	<input type="text" value="MyCompany.ad"/>
Kerberos realm.	<input type="text" value="MYCOMPANY.AD"/>
The default group an Kerberos authenticated user gets when there is no explicit group assigned.	<input type="text" value="EVERYBODY"/>
<input type="button" value="Save Kerberos Settings"/> <input type="button" value="Test"/> <input type="button" value="Refresh"/>	

Figure 52 - Settings for configuring OO with Kerberos authentication

To authenticate with Kerberos

1. Select the **Kerberos Enabled** check box.
2. In the **Value** box for the **Kerberos 5 configuration file**, type the name of the file.
The location of the file should be within the OO home directory, and the path should be relative to that directory.

For instance, the example in the description in the **Description** of the setting describes the Kerberos configuration file as being in the \Central\conf subdirectory of the OO home directory.

Important! If the specified file path for the Kerberos configuration file is invalid, OO authenticates users based on the platform's default Kerberos configuration file, if any. (Usually, on a Windows system, this is C:\Windows\krb5.ini, and on a Linux system, it is /etc/krb5.conf.) As a result of this default behavior, unintended authentications can take place. Note that there is no error message to alert you to this condition.

3. In the **Value** box for the **KDC host**, type the IP address or fully qualified machine name of the Key Distribution Center (KDC), the authentication center for users.

Use the following syntax:

```
<KDC_host>[:<port>]
```

where:

- `<KDC_host>` is the IP address or fully qualified name of the LDAP server.
- `<port>` is the port number that the KDC host uses, if the KDC host is configured to use a non-default port.

For instance, if the LDAP server is mirage.ad, its IP address is 192.168.5.5, and it uses port 200, the setting would be:

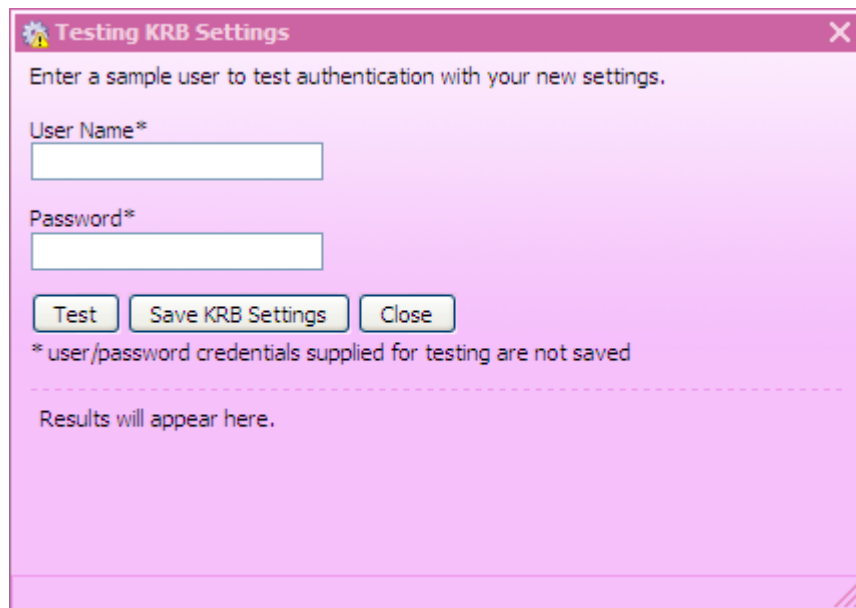
```
mirage.ad:200
```

or

```
192.168.5.5:200
```

4. In the **Value** box for the **Kerberos realm**, type the domain name of the realm.
For instance, the domain might be **MIRAGE.AD**.
5. In the **Value** box for the **default group**, specify which OO group or individual is mapped to when a mapping is not specified.
For instance, to assign unmapped groups or individuals only the capability to run flows, you would type **LEVEL_ONE**. To learn about the default OO groups, see [Managing groups](#).
6. To save your settings, click **Save Kerberos Settings**.
7. To test the current Kerberos authentication settings from within OO, click the **Test** button.

The **Testing KRB Settings** dialog box appears.



8. Type the user name and password of an external account that is authenticated by Kerberos, and then click **Test**.
9. If the test fails, modify the settings and test again.
10. When the authentication settings test successfully, click the **Save KRB Settings** button.
11. When you see a message to restart the Central service (RSCentral on a Windows system or Central.sh on a Linux system), do so.

If your IT organization also authenticates with Active Directory and/or LDAP, complete the procedures in [AD authentication settings or LDAP authentication settings](#).

Finally, because Kerberos authenticates only users, you manually assign external Kerberos users (rather than their groups) to OO groups. For more information, see [Managing users](#).

Managing users

Users are either **internal** to OO—that is, you create them within OO, which they do not exist outside of—or **external** user accounts that exist independently of OO, such as Active Directory or LDAP accounts.

- When you create an internal user, you also create a password and assign the user to one or more OO groups.
- When you add an external user, you do not create the account’s password, and you must either assign the account to one or more OO groups, or map the user’s external (AD or LDAP) account to an OO account.

Necessary capabilities: To manage user accounts, you must be a member of a group that has the `MANAGE_USERS` capability.

Internal or external users?

The rule of thumb is that you create internal accounts for use in testing environments, which may be isolated from the domains or directories through which external users are authenticated, and that, where Central is installed in a production environment, you add users from external domains or directories (and map their groups to OO groups). Adding external users is less work than creating internal users.

For example, suppose you have a staging Central server in a testing environment and a production Central server.

- The staging server might have only two or three flow authors as users, so it would make sense to create internal OO users for those authors to log on to Central with when testing their flows.
- The production server, however, might have two dozen or so IT personnel logging on to Central in order to run flows, in addition to administrators and managers who might need to log on in order to create charts for analyzing the data generated by the flows. In this case, you would probably want to add external users and map their external groups to OO groups.

Note: If an external group has the same name as an internal OO group (after translation of the external name using OO group-name rules), then members of the external group can log on to OO with the capabilities of the OO group. OO group-name rules are that names are all upper-case and spaces are replaced by underscores.

Thus, if “htudor” belonged to an external group named Level One, then “htudor” would be able to log in to OO with the capabilities of the OO group LEVEL_ONE. However, if the name of “htudor’s” external group were Level 1, “htudor” would not be able to log in to OO.

Adding a user

Necessary capabilities: To add a user and add the user to an OO group, you must be a member of a group that has, respectively, the MANAGE_USERS and the MANAGE_GROUPS capabilities.

To add an individual user

1. On the **Administration** tab’s **Manage Users** subtab, under **Users**, click **Add New User**. The **User Information** dialog appears.




Figure 53 - User Information dialog box

2. Type the user account name.
3. If you are creating the account within OO, select the **Internal Account** check box, and then type and verify a password with which the user can log on to OO.

Note: You can use this feature to add a specific internal user to an OO group, for example in a staging environment.

The password must be at least six characters long.

Note that by default, the **Account Enabled** check box is selected.

4. To assign the new user to a group, click the **Assigned Groups** tab within the dialog and then select the check boxes for the groups whose capabilities the user should have. For information on the capabilities that are assigned to the various groups, see the *OO Administrator's Guide* (AdminGuide.pdf).

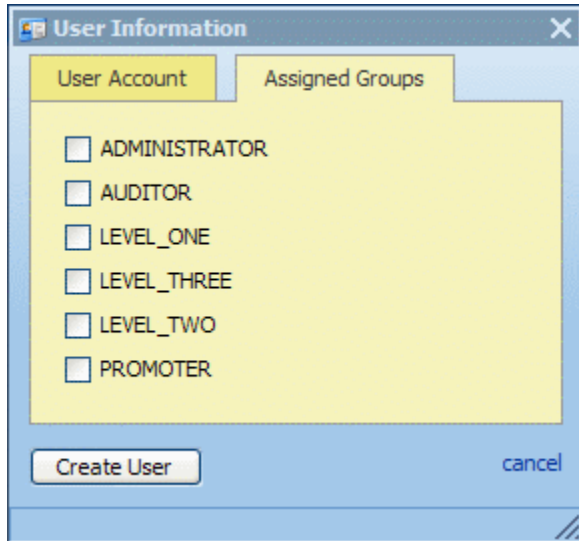


Figure 54 - Assigning groups to a user

5. To finish, click **Create User**.

Editing a user's account

Note that you cannot edit the admin account. The admin account possesses all capabilities, and neither can nor should be altered. As long as you have the admin account password, it provides access to all parts of OO even if all other accounts are disabled or otherwise nonfunctional.

Necessary capabilities: To modify a user, you must be a member of a group that has the `MANAGE_USERS` capability. If your changes include adding the user to or removing him from an OO group, your group must also have the `MANAGE_GROUPS` capability.

To make changes to a user's account

1. On the **Administration** tab's **Manage Users** subtab, click the Edit icon (📄) next to the name of the user whose account you want to change, then, in the **User Information** dialog, make changes in the same way that you configured the account when you added the user.
For more information, see [Adding a user](#).
2. When you've finished making changes, click **Update User**.

Deleting a user

Necessary capabilities: To delete a user, you must be a member of a group that has the `MANAGE_USERS` capability.

To delete a user account

- On the **Manage Users** subtab of the **Administration** tab, select the check box for the user in the **Delete** column, and then click **Delete Selected**.

Managing groups

Groups are the basic unit for defining users' scope of activity. You exercise this control by assigning groups:

- Capabilities, or attributes that determine which actions the members of a group can do.
Note: There is not a capability for executing flows. Access to each flow for running it is controlled by its author, who selectively grants the Execute access permission for the flow when he or she creates it.
- Access permissions, which determine which flows, operations (and other parts of flows, such as domain terms) that members of a group can work on.

Necessary capabilities: To make changes to groups, you must be a member of a group that has the MANAGE_GROUPS capability.

OO groups are integral to adding users as groups to OO and to managing their capabilities. You can map your IT organization's AD or LDAP groups to OO groups, thus adding the entire membership of the group at once, as OO users.

Scenario

Suppose you want to map the following groups in your IT organization to existing OO groups.

Service Desk	Network Specialists	Managers
These are front-line Help desk IT personnel. They need to be able to start and, probably, schedule flows. To let them run and schedule but not author flows in OO, you could empower them to view and analyze data generated by the flows.	Members of this group have the expertise to author flows. They will run them at least as part of testing.	Let's suppose that these users need to harvest and analyze information from flows, but they don't need to start or author flows.

Before we look at how we might map these groups to OO groups and create rights for the OO groups that make sense for these external groups, let's look at the groups that are created by default when you install Central:

- The LEVEL_ONE, LEVEL_TWO, and LEVEL_THREE groups are user groups whose rights you define with capabilities and access permissions.
 You cannot delete the remaining four default OO groups (ADMINISTRATOR, AUDITOR, EVERYBODY, and PROMOTER), which have different specific purposes.
- ADMINISTRATOR
 The purpose of the ADMINISTRATOR group is to have one account that you can use to run and work in Central and Studio in case you temporarily lose the ability to log users in. This group possesses all capabilities and access permissions, so you cannot modify its capabilities or access permissions. You can, however, change the password.
 Although you can add members to this group, keep in mind that it is an all-powerful group within OO, so you should assign this account to the fewest people possible.
- AUDITOR
 As the description indicates, the AUDITOR group might be appropriate for administrators and managers, who should be able to see the data that flows have generated, but who should not necessarily run or author flows. Members of this group have Read permission on all objects and have capabilities that allow them to view flow schedules and create reports.
- EVERYBODY

Every user that you add to OO automatically becomes a member of this group. The group doesn't have any capabilities, but does have access to certain OO objects, such as Accelerator Packs. As a result, the OO administrator's maintenance tasks are reduced. Further, this group's existence enables authors to give Read, Write, or Execute permission for a flow to everyone at once, if desired, instead of having to grant access permissions group by group.

- **PROMOTER**

The PROMOTER group has no capabilities, but is able to publish to any Central repository.

This group is intended to limit those who can publish to the various Central (public) repositories. For installations of Central in production environments (and possibly also in staging environments), as opposed to development environments, restricting the membership of this group to just one person or a few people helps ensure that only work that is ready for the staging or production environment is published to it.

The following schematic shows the promotion of flows.




Figure 55 - Promotion of Central repositories

Thus authors who are not members of the PROMOTER group can work freely in the public Central repository in the development environment. When a flow is ready to be tested in the staging/testing environment, the PROMOTER publishes the flow from the development Central repository to the staging Central repository. Then, once the flow has been sufficiently tested, the PROMOTER for the production environment publishes the flow from the staging Central repository to the production Central repository. Having just one person use the promote repositories reduces the danger of a flow that is not ready for the testing or production environment being published there before it is ready.

Default groups on the Manage Groups subtab

The following shows the default groups as they are listed and described on the **Manage Groups** subtab of the **Administration** tab. The groups that cannot be deleted are represented in red.










Edit	Group Name ▲	Description	Capabilities	External Groups Mapping	Delete [Check all <input type="checkbox"/>]
	ADMINISTRATOR	Represents Operations Orchestration administrators.	MANAGE_USERS, MANAGE_GROUPS, AUTHOR, SCHEDULE, MANAGE_RUNS, RUN_REPORTS, MANAGE_CONF, VIEW_SCHEDULES, HEADLESS_FLOWS		
	AUDITOR	Represents Operations Orchestration auditors. Users from this group have unconditional read access to the repository.	RUN_REPORTS, VIEW_SCHEDULES		
	EVERYBODY	Every authenticated user is part of this group.	NONE		
	LEVEL_ONE	Represents Operations Orchestration level one users.	NONE		 <input type="checkbox"/>
	LEVEL_THREE	Represents Operations Orchestration level three users.	NONE		 <input type="checkbox"/>
	LEVEL_TWO	Represents Operations Orchestration level two users.	NONE		 <input type="checkbox"/>
	PROMOTER	Represents Operations Orchestration promoters. Users from this group have unconditional read access to the repository as well the ability to publish changes to any repository.	NONE		

Figure 56 - Manage Groups subtab

You use the **Manage Groups** subtab of the **Administration** tab to:

- Map external groups to a group.
- Change a group's capabilities.
- Change the group's description or name.

Adding groups

Necessary capabilities: To create a group, you must be a member of a group that has the MANAGE_GROUPS capability.

To add a group

1. On the **Administration** tab, click **Manage Groups**.
2. In the **Groups** box, click Add **New Group**.
The **Group Information** dialog box appears.




Figure 57 - Creating a group: Group Information

3. Type a name for the group and, if you wish, a description of the group.
4. To specify the group's capabilities:
 - a. Click the **Assigned Capabilities** tab of the **Group Information** dialog box.
 - b. Select the capabilities that the group needs, and click **Create Group**.

For more information on the role of group capabilities and on assigning them, see [Managing groups](#) and [Changing a group's assigned capabilities and description](#).

5. To map an external group to the group you have created:
 - a. Click the **External Groups Mapping** tab of the **Group Information** dialog box.
 - b. Type a comma-separated list of the external group names that you want to map to this group, and click **Update Group**.

For more information on mapping external groups to internal OO groups, see [Mapping external groups to OO groups](#). For more information on using external groups, see the *OO Administrator's Guide* (AdminGuide.pdf).

Adding OO users to groups

After you have created an OO user (on the **Administration\Manage Users** subtab), you assign the user to one or more OO groups.

Necessary capabilities: To add a user to a group, you must be a member of a group that has the `MANAGE_USERS` capability.

To assign a user to an OO group

1. On the **Administration** tab, click the **Manage Users** subtab.
2. In the row for the user you want to add to a group, click the **Edit** icon.
3. In the **User Information** dialog box, click the **Assigned Groups** tab, and then specify the groups you want to add the user to.

Mapping external groups to OO groups

To add users of the following AD or LDAP groups to OO groups as shown in the following table, use the **Manage Groups** subtab on the **Administration** tab. For example, if your LDAP defines the (external to OO) groups “Service Desk,” “Network Specialist,” and “Quality Assurance,” you might map them to OO groups as in the following table.

AD or LDAP group	OO groups
Service Desk	LEVEL_ONE
Network Specialist	LEVEL_THREE
Quality Assurance	LEVEL_TWO, PROMOTER

Notes:

- To map external groups to OO internal groups, you must enable an external authentication provider. To do so, see [Using external authentication for Central users](#) and one of the following sections:

- [AD authentication settings](#)
- [LDAP authentication settings](#)
- [Kerberos authentication settings](#)

To enable an external authentication provider, you must be a member of a group that has the MANAGE_CONF capability.

- In this example, you can define what the LEVEL_ONE, LEVEL_TWO, and LEVEL_THREE groups can do by assigning capabilities to the group. So to enable the LDAP group Quality Assurance to test flows, you might assign the LEVEL_TWO group in OO the SCHEDULE, MANAGE_RUNS, VIEW_SCHEDULES, and HEADLESS_FLOWS capabilities.

The PROMOTER group is a default group, and you cannot change its capabilities.


Authors define which objects a group can act on by changing the group permissions for the object. For more information on capabilities and permissions, see [Capabilities and access permissions](#).

- If an external group has the same name as an internal OO group (after translation of the external name using OO group-name rules), then members of the external group can log in to OO with the capabilities of the OO group. OO group-name rules are that names are all upper-case and spaces are replaced by underscores.

Thus, if htudor belonged to an external group named Level One, then htudor would be able to log in to OO with the capabilities of the OO group LEVEL_ONE. However, if the name of htudor’s external group were Level 1, htudor would not be able to log in to OO.

Necessary capabilities: To map an external group to an OO group, you must be a member of a group that has the MANAGE_GROUPS capability.

To map external groups to OO groups

1. On the **Administration** tab, click the **Manage Groups** subtab, and then click the **Edit** icon () beside the group you want to map the external group to.
2. In the **Group Information** dialog that appears, click the **External Groups Mapping** tab.

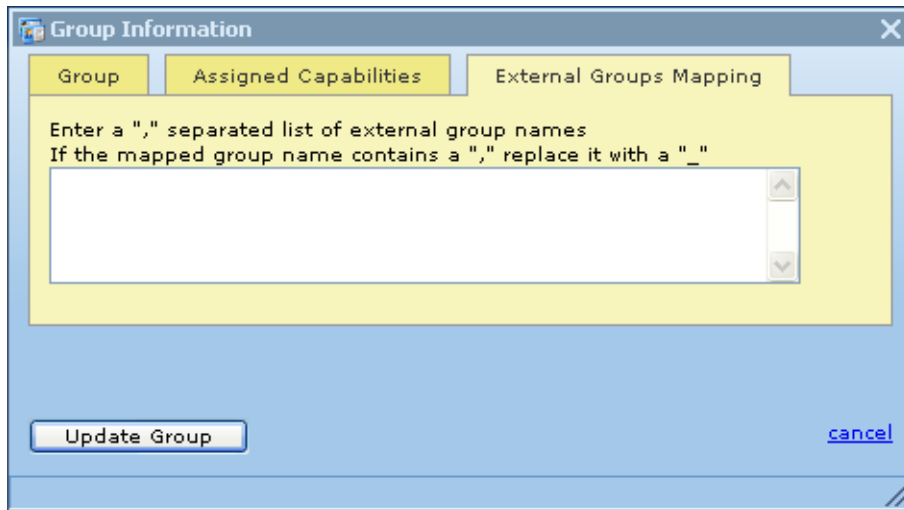


Figure 58 - External Groups Mapping tab

3. In the text box, type the name(s) of the external group or groups whose members you want to be members of this OO group.
For instance, to map the AD or LDAP "Network Specialist" group to the OO LEVEL_THREE group, type **Network Specialist** in the text box.
4. Click **Update Group**.

You create and manage OO user accounts, manage group membership, and assign capabilities (defined actions) on the **Administration** tab of Central. For more on OO groups and capabilities, see the *OO Administrator's Guide*; for the procedure for assigning a capability to a group or user, see [Changing a group's assigned capabilities and description](#), below.

Permissions are granted by flow authors in Studio. The available permissions are Read, Write, Execute, and Link to permissions to flows and objects that are associated with them. For more information on granting permissions, see Help for Studio.

Changing a group's assigned capabilities and description

Capabilities are the actions that OO enables members of a group to perform within the OO system of components. Capabilities are half of the OO system for controlling access to flows and other OO repository objects. The other half is permissions for the object. A user or author must have the appropriate permissions to run or schedule a flow. For more information on OO repository objects and permissions, see Help for Studio or the *Guide to Authoring Operations Orchestration Flows* (Studio_AuthorsGuide.pdf).

Necessary capabilities: To modify a group, your group must have the MANAGE_GROUPS capability.

To change the capabilities assigned to a group

1. On the **Administration** tab, click the **Manage Groups** subtab, and then click the Edit icon (📄) in the row of the group whose capabilities and description you want to change.
2. In the **Group Information** dialog, click the **Assigned Capabilities** tab.




Figure 59 - Assigning capabilities for a group

3. Select the check boxes for the capabilities that you want the members of this group to have.

Notes:

- You cannot change the capabilities for the ADMINISTRATOR group or AUDITOR group. For information on the intended uses of these groups, see the *Operations Orchestration Administrator's Guide* (AdminGuide.pdf).
 - By default, the groups LEVEL_ONE, LEVEL_TWO, and LEVEL_THREE have no capabilities, so you must assign them some.
 - For information on capabilities and the difference between them and access permissions to objects, see the *Operations Orchestration Administrator's Guide*.
4. If your conception of the group changes after you change the capabilities you grant it, you may want to click the **Group** tab and change the group's name and description to make them more descriptive.
 5. When you've finished working here, click **Update Group**.

In addition, the following tasks must be executed outside the Central application. For more information on these tasks, see *Operations Orchestration Administrator's Guide*.

- Configuring Active Directory to run over SSL.
- Configuring OO for extended functionality (with Java Remote Action Service and .NET Remote Action Service).
- Changing the Studio configuration in the Studio.properties file.
- Backing up OO, including all Studio repositories and the Central database of run-history information.

Deleting groups

You cannot delete the ADMINISTRATOR, AUDITOR, or EVERYBODY groups. For more information on these groups, see [Managing groups](#).

Necessary capabilities: To modify a group, your group must have the MANAGE_GROUPS capability.

To delete a group

1. On the **Administration** tab, click **Manage Groups**.

The **Manage Groups** subtab appears as follows.










Edit	Group Name ▲	Description	Capabilities	External Groups Mapping	Delete [Check all <input type="checkbox"/>]
	ADMINISTRATOR	Represents Operations Orchestration administrators.	MANAGE_USERS, MANAGE_GROUPS, AUTHOR, SCHEDULE, MANAGE_RUNS, RUN_REPORTS, MANAGE_CONF, VIEW_SCHEDULES, HEADLESS_FLOWS		
	AUDITOR	Represents Operations Orchestration auditors. Users from this group have unconditional read access to the repository.	RUN_REPORTS, VIEW_SCHEDULES		
	EVERYBODY	Every authenticated user is part of this group.	NONE		
	LEVEL_ONE	Represents Operations Orchestration level one users.	NONE		 <input type="checkbox"/>
	LEVEL_THREE	Represents Operations Orchestration level three users.	NONE		 <input type="checkbox"/>
	LEVEL_TWO	Represents Operations Orchestration level two users.	NONE		 <input type="checkbox"/>
	PROMOTER	Represents Operations Orchestration promoters. Users from this group have unconditional read access to the repository as well the ability to publish changes to any repository.	NONE		

Figure 60 - Managing groups

2. On the **Manage Groups** subtab, in the **Delete** column, in the row for the group that you want to delete, either click the red button or select the box and then click **Delete Selected**.

OR

To delete all the groups (except those that cannot be deleted), in the header for the **Delete** column, select the **Check all** box, and then click **Delete Selected**.

Managing flow runs

On the **Current Runs** tab, you can view current flow runs and resume, delete, or reassign them.

The main tasks in managing runs are the following. Procedures for these tasks are described below:

- Viewing current runs
- Deleting runs
- Reassigning runs
- Resuming runs

Necessary capabilities: To manage flow runs, your group must have the MANAGE_RUNS capability.

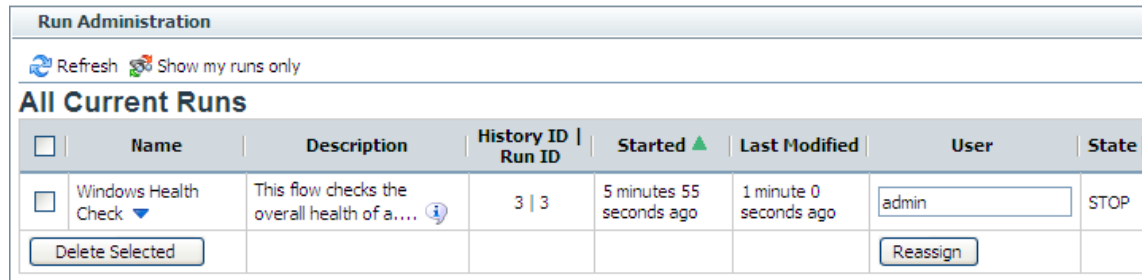
To see which flows are currently running

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Current Runs** navigation tab.
3. To show only the runs that you have ownership of (usually because you started the runs), click **Show my runs only**.

OR

To show all the current runs regardless of who owns them, click **Show all users' runs**.

The **Run Administration** table displays current flow runs, including the following information.



The screenshot shows the 'Run Administration' interface. At the top, there are 'Refresh' and 'Show my runs only' buttons. Below this is a section titled 'All Current Runs' containing a table with the following columns: Name, Description, History ID | Run ID, Started, Last Modified, User, and State. One row is visible for 'Windows Health Check' with a description 'This flow checks the overall health of a...', '3 | 3' in the History ID column, '5 minutes 55 seconds ago' in the Started column, '1 minute 0 seconds ago' in the Last Modified column, 'admin' in the User column, and 'STOP' in the State column. Below the table are 'Delete Selected' and 'Reassign' buttons.

<input type="checkbox"/>	Name	Description	History ID Run ID	Started ▲	Last Modified	User	State
<input type="checkbox"/>	Windows Health Check ▼	This flow checks the overall health of a... ⓘ	3 3	5 minutes 55 seconds ago	1 minute 0 seconds ago	admin	STOP

Figure 61 - Administering Runs

Note: The page does not refresh by itself; you must refresh it to reflect any runs that have been started, interrupted, handed off, or resumed.

4. To refresh the page, click **Refresh**.

You can resume or delete a run or view the run's history by clicking the downward pointing arrow beside the name of the run. For more information:

- To resume a run, see [Resuming a run](#).
- To delete a run, see [Deleting a run](#).
- To view a run's history, see [Run histories: What happened and why](#).

Other system configurations

If you are a member of a group that has the `MANAGE_CONF` capability, you can customize several other aspects of the OO system:

- [Whether Return on Investment \(ROI\) reporting is visible](#)
- [Whether only administrators can author in the Central repository](#)
- [The maximum number of versions of an object are kept in the repository](#)
- [Whether only members of PROMOTER or ADMINISTRATOR group can publish](#)
- How frequently the Dashboard charts refresh
- What prefix, if any, for flow input names should be required for any flow that you use a URL to start a run of in the Central Web page.
- For information on specifying a prefix that must be used for all flow input names used in URLs, see the section below, [Specifying a required prefix for init params when using a URL to start a flow run in Central](#). Additional information is in the *OO SDK Guide* (SDKGuide.pdf).
- [How LDAP referrals from the primary LDAP server to other LDAP servers are handled](#)
- [Whether to automatically resume headless runs that were interrupted by a Central server's failure](#)
- [Configuration of a Central server cluster and whether it is enabled](#)

You make these changes on the Central web page's **Administration** tab, on that tab's **System Configuration** subtab.

Dashboard | Flow Library | Current Runs | Reports | Scheduler | **Administration**

Account | Manage Users | Manage Groups | Dashboard Chart Definitions | **System Configuration**

▼ Central Settings

Refresh

General Settings

Description	Value
When set to true, only users in the ADMINISTRATOR group can see ROI (Return On Investment) in reports.	<input type="text" value="true"/>
When set to true, only users in the ADMINISTRATOR group may author in the repository. Recommended for a production environment.	<input type="text" value="false"/>
Maximum number of versions for an object that will be kept in a repository.	<input type="text" value="50"/>
When set to true, only users in the PROMOTER or ADMINISTRATOR group may publish. Recommended for a production environment.	<input type="text" value="false"/>
Time interval measured in minutes representing the refresh rate for the aggregate data on the dashboard page. If not set, it defaults to 1 minute.	<input type="text" value="1"/>
Prefix for init params for flow invocation through URLs using the GUI. Init parameters names *MUST* always be different from "service" or "sp". When this is defined all init parameters *MUST* start with the specified prefix.	<input type="text"/>
How to handle LDAP referrals. Valid values are "follow", "ignore", "throw". Used only if AD or standard LDAP are enabled.	<input type="text" value="follow"/>
Automatically resume orphaned headless runs from a Central failure.	<input type="text" value="false"/>

Figure 62 - Administration tab, System Configuration subtab

Enabling the visibility of ROI reporting

By default, ROI is only visible to users in the ADMINISTRATOR group in:

- A column in the reports created on the **Reports** tab.
- The **Flow ROI Value** column in the **Popular Flows** chart on the **Dashboard**.

However, you can make ROI visible to all users.

Necessary capabilities: To change the configuration of OO, your group must have the MANAGE_CONF capability.

To enable or disable ROI reporting visibility

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. In the **General Settings** area:
 - To allow all users to see ROI reporting, in the **When set to true, only users in the ADMINISTRATOR group can see ROI** row, type **false** in the **Value** box.
 - OR
 - To allow only users in the ADMINISTRATOR group to see ROI reporting, in the **When set to true, only users in the ADMINISTRATOR group can see ROI** row, type **true** in the **Value** box. (This is the default setting.)
4. Click **Save General Settings**.

5. Restart the OO Central service (RSCentral).

Restricting authoring to administrative users

In a production environment, you may want to restrict the ability to change flows and other OO objects to very few people, even to a single person. Doing so helps prevent unwanted changes from appearing in flows that are used in your “real world” outside the development/staging environment. One way to do so is to configure your OO system so that only members of the ADMINISTRATOR group can perform authoring tasks in the Public repository in your production environment.

Necessary capabilities: To change the configuration of OO, your group must have the MANAGE_CONF capability.

To restrict authoring to administrative users

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. To restrict the authoring capability to members of the ADMINISTRATOR group, in the **General Settings** area, in the **When set to true, only users in the ADMINISTRATOR group may author in the repository** row, type **true** in the **Value** box. The default value is **false**.
Note: If you have upgraded to OO version 7.60 from an earlier version, this value will default to **true** unless the default was changed to **false** in the previous version.
OR
To allow other users who are members of groups that have the AUTHOR capability, leave **false** in the **Value** box.
3. Click Save General Settings.
4. Restart the OO Central service (RSCentral).

Setting how many versions of an object are stored

A version of an OO object is created when an author saves his or her changes to the object. With the Show History command in Studio, authors can recover earlier versions of a given OO object. To prevent the Central repository from growing too much, you can set a maximum for the number of versions of an OO object that are stored. After this maximum is reached, as each new version is created by a checkin, the oldest version is permanently deleted.

Necessary capabilities: To change the configuration of OO, your group must have the MANAGE_CONF capability.

To specify how many versions of an OO object are stored

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. In the **General Settings** area, in the **Maximum number of versions for an object** row, in the **Value** box, type the number of versions that you want stored.
 - The number 0 is a valid value, so typing **0** specifies that no versions are saved.
 - If you leave the **Value** column empty or type an invalid value, Central uses the default value of **50**.
4. Click **Save General Settings**.
5. Restart the OO Central service (RSCentral).

Restricting who can publish

Allowing only a small group of people, or even just one person, to publish to the Central repository in your production environment can help prevent undesired versions of Ops flows, such as those that have not been sufficiently tested, from being used where any errors in the flows could affect your organization negatively.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

To restrict the ability to publish to members of the **ADMINISTRATOR** or **PROMOTER** groups

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. In the **General Settings** area, in the **When set to true, only users in the PROMOTER or ADMINISTRATOR group may publish** row, in the **Value** box, type **true**. The default value is **false**.
Note: If you have upgraded to OO version 7.60 from an earlier version, this value will default to **true** unless the default was changed to **false** in the previous version.
4. Click **Save General Settings**.
5. Restart the OO Central service (RSCentral).

Changing the Dashboard charts refresh rate

On the **Administration** tab in Central, you can change how frequently Central Dashboard charts are updated.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

To change the rate at which Dashboard charts refresh their data

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. To change how frequently Dashboard charts are updated with new data, in the **General Settings** area, in the **Time interval...refresh rate** row, in the **Value** box, type a whole number reflecting the number of minutes you want between updates. The default value is **1**.
4. Click **Save General Settings**.
5. Restart the OO Central service (RSCentral).

Specifying a required prefix for init params when using a URL to start a flow run in Central

Flow and step inputs are treated as initial parameters (also known as “init params”) when you start a flow in the Central Web page with a URL. The initial parameters input names “sp” or “service” are both reserved, and using them when you use a URL to start a flow in the Central Web page can cause problems. To prevent the use of inputs named “sp” or “service”, you can specify a required prefix for all init params used in a URL that starts a flow in Central.



Key information: Important! The prefix that you specify is required only for init params in the URL that starts the flow run in Central. This prefix is not relevant to input names specified for flows in Studio. It is, however, required for all init params specified in URLs that start flows in

the Central Web page.

With a prefix defined, even if a flow has inputs named “sp” or “service”, defining the required prefix and then prepending the prefix to the input names as init params in a URL that starts the flow prevents problems that would occur otherwise.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

For more information, including the procedure for specifying a required prefix and characters to avoid using in the required prefix, see the *OO SDK Guide* (SDKGuide.pdf).

Specifying how Central manages LDAP referrals

When you have enabled Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) authentication, you can specify how referrals from what Central authentication does when it encounters an LDAP referral from one server or namespace to another. You can specify that Central do one of the following:

- Follow the referral.
- Ignore (that is, not follow the referral).
- Throw an exception.

Note: This is relevant only if you have enabled AD or LDAP authentication.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

To specify how Central authentication manages LDAP referrals

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. On the **How to handle LDAP referrals** line, in the **Value** box, type either **follow**, **ignore**, or **throw**, depending on how you want Central authentication to respond. The default setting is **follow**.
4. Click **Save General Settings**.
5. Restart the OO Central service (RSCentral).

Enabling Central to resume runs that were interrupted when Central failed

Headless flow runs are runs that have not been manually started from within Central. No Central user is considered to be the owner of such a run. If a headless run is halted by a failure of Central, its lack of an owner means that there is no user to restart it once Central is running once again. In this case, the headless run is considered to be orphaned. To remedy this situation, you can configure Central to automatically resume orphaned headless runs. This configuration applies to nonclustered Central installations as well as clustered ones.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

To enable Central to resume orphaned headless runs after Central recovers from failure

1. Log on to Central with an account that has OO administrative permissions.
2. On the Central **Administration** tab, click the **System Configurations** subtab.
3. Under **General Settings**, in the **Value** box for the **Automatically resume orphaned headless runs** setting, replace **false** with **true**.

4. Click **Save General Settings**.
5. Restart the OO Central service (RSCentral).

Changing the configuration of a Central cluster

The only settings for failover and run recovery that you change in Central are:

- The name of the cluster that the Central server is a member of.
- Whether clustering is enabled.

Necessary capabilities: To change the configuration of OO, your group must have the `MANAGE_CONF` capability.

For more information on Central failover and run-recovery clusters, see *Installing or Upgrading HP Operations Orchestration* (InstallGuide.pdf).

To enable clustering and select a cluster for the Central server to join

1. Click the **Administration** tab, and then click that tab's **System Configuration** subtab.
2. Scroll down to the **Clustering Settings** section.

Clustering Settings	
Description	Value
The name of the cluster this Central belongs to.	CENTRAL_CLUSTER
<input type="button" value="Save Clustering Settings"/> <input type="button" value="Refresh"/>	

Figure 63 – Joining the Central server to a cluster

3. In the **The name of the cluster** row, in the **Value** box, type the name of the cluster that has been created for Central.
4. Click **Save Clustering Settings**.

For information on changing the network address binding on Central servers that have more than one network interface, see *Installing or Upgrading HP Operations Orchestration* (InstallGuide.pdf) or the *OO Administrator's Guide* (AdminGuide.pdf).

Troubleshooting

Web browser shows a security-certificate-related warning when you open the Central Web site

You can safely proceed past this warning.

For installations of Central that communicate using the HTTPS protocol, Web browsers show security violation errors or messages unless your Web administrator creates a valid security certificate for delivering the Central Web pages. If you see such a browser warning, it is because OO includes, by default, an unsigned certificate that serves as a placeholder for a valid customer-obtained certificate. If you choose not to create a security certificate, you can safely ignore the warning.

I was sent back to the login page

Your login may have timed out. Log in again.

I cannot restart or resume a flow that I started or that was handed off to me

If the Central service on the Central server fails or is stopped, any runs that were active are left without owners. This means that if an active run that you started is interrupted because the Central service (RSCentral) stopped, the Central user who had ownership of the flow (the user who started it or to whom it had been handed off) is no longer the owner, and must be a member of the ADMINISTRATOR group to restart it.

I cannot change or create a schedule for a flow

Check with the flow author to see whether the groups that you are a member of have Write permission for the flow you're trying to schedule.

Make sure that one of the groups that you are a member of has the VIEW_SCHEDULES and SCHEDULE capabilities.

Central fails to create a schedule for a flow

1. Check your database connection information to ensure that Central is connected to the database.
2. Restart the RSScheduler service.
3. Refresh the browser page for Central.

Changes were made to the Public repository, but they don't appear in the Flow Library

Changes to Central's repository do not appear in the Flow Library in Central until the Flow Library has been reloaded. To reload the Library, click the **Flow Library** tab.

Flows run under heavy load with command-line or RAS operations fail on Windows systems with error code 128

If the Central server or a standalone RAS installation has a Windows operating system, flows that run on those machines under intense usage and using command-line operations or the RAS operation may fail with error code 128. This can result from Data Execution Prevention (DEP) being enabled for all programs. By default, DEP is enabled for all programs and services except any exceptions that you specify. Alternatively, try enabling DEP "for essential Windows programs and services only." For information on changing the DEP settings, see Help for Windows.

Index

- Access control
 - and groups, 49
 - and users, 49
- Access to Central for external users, 52
- AD authentication
 - configuring, 53
 - enabling, 53
- Administration tab
 - Manage Groups subtab, 65
- Administrative users
 - restricting authoring to, 74
- ADMINISTRATOR group
 - defined, 64
 - restricting authoring to, 74
- AUDITOR group
 - defined, 64
- Authentication. *See* Kerberos authentication, *See* LDAP authentication, *See* AD authentication
- Authoring
 - restricting to administrative users, 74
- Capabilities, 50
 - defined, 50
- Central
 - flow Metrics area, 3
 - navigating in, 5
 - navigation tabs, 5
 - Popular flows, 4
 - starting, 2
 - troubleshooting, 77
- Central authentication
 - and LDAP referrals, 76
- Central clusters
 - reconfiguring, 77
- Central users
 - and external authentication, 53
- Charts, reporting, 37, 38
- CI. *See* Configuration Item
- Configuration Item
 - defined, 36
- copyright notices, ii
- Current Runs tab, 18
- Current user
 - changing password, 5
- Dashboard. *See also* Reporting charts
 - overview, 35
 - refresh rate, changing, 75
 - reporting charts, 35
- Dashboard charts
 - importing, 42
 - interpreting, 37
- EVERYBODY group
 - defined, 64
- External authentication
 - and Central users, 53
- External groups
 - mapping to OO groups, 68
- External users
 - providing access to Central, 52
- Flow graph
 - opening in a new window, 16
- Flow Library tab, 6
- flow Metrics area, 3
- flow Metrics graph
 - customizing, 4
- Flow previews
 - navigating in, 11
- Flow run
 - specifying required prefix for URL, 75
- flow runs
 - bookmarking, 46
 - creating a link to, 46
 - deleting, 44
 - handing off, 48
 - interrupting, 44
 - reassigning ownership, 44
 - resuming, 44, 45
 - resuming when handed off, 48
- Flow runs. *See* flow runs
 - deleting, 45
 - managing, 71
 - monitoring, 18
- Flow schedules. *See* Schedules
- Flows. *See* flows
 - auditing, 49

- browsing, 7
 - choosing, 6
 - creating links to, 47
 - defined, 1
 - finding, 6, 8
 - guided run, 12
 - instant run, 12
 - managing, 49
 - previewing, 9
 - run all, 12
 - running, 12
 - scheduling, 28
 - Scheduling simultaneous runs, 29
 - searching for, 8
 - starting from outside Central, 43
 - viewing, 9
- Groups
- adding, 66
 - adding users, 67
 - and access control, 49
 - capabilities, 50, 69
 - default, 64, 65
 - defined, 64
 - deleting, 70
 - descriptions, 69
 - managing, 64
 - mapping external groups to, 68
 - permissions, 51
 - viewing, 6
- Headless flow runs
- automatically restarting when orphaned, 76
- Histories, run. *See* Run histories
- Individuals
- permissions, 51
- Kerberos authentication
- configuring, 59
 - enabling, 59
- LDAP authentication
- configuring, 57
 - enabling, 57
- LDAP referrals
- Central responses to, 76
- legal notices, ii
- copyright, ii
 - restricted rights, ii
 - trademark, ii
 - warranty, ii
- LEVEL_ONE group
- defined, 64
- LEVEL_THREE group
- defined, 64
- LEVEL_TWO group
- defined, 64
- Manage Groups subtab, 65
- My Current Runs, 18
- Objects
- permissions, 51
- OO Central
- Quick Overview, 1
- OO groups. *See* Groups
- OO users. *See* Users
- Operations Orchestration Central. *See* OO Central
- Password
- changing, 5
- Permissions, 51
- defined, 50
- Popular flows, 4
- PROMOTER group
- defined, 64
- Publishing
- restricting rights to, 75
- Reporting charts, 35
- creating, 39
 - modifying, 39
- restricted rights legend, ii
- Return on Investment. *See* ROI
- ROI
- disabling reporting, 73
 - enabling reporting, 73
- Rflowinvoke.exe, 43
- Run histories
- more detail, 25
 - overview, 21
 - selecting columns, 25
 - viewing, 21
 - viewing step results, 27
- Run reports
- creating, 21
 - specifying, 21
 - viewing, 21
- Runs. *See* flow runs
- Scheduler
- changing how it works, 33
 - configuring, 33
- Schedules
- creating, 29
 - deleting, 33
 - disabling, 33

- editing existing, 33
- enabling, 33
- Scheduler, configuring, 33
- working with, 32

Starting flow run with URL

- specifying required prefix for, 75

Subflows

- running, 16

System configurations, 72

Tabs

- Current Runs, 18
- Flow Library, 6
- Reports, 21
- Scheduler, 28

trademark notices, ii

Troubleshooting, 77

User accounts. *See* Users

User groups. *See* Groups

Users

- adding, 62
- adding to groups, 67
- and access control, 49
- deleting, 63
- internal vs. external, 61
- managing, 61
- user account, editing, 63

warranty, ii