HP Operations Orchestration Software Central

Software Version: 7.20

Users' Guide

Document Release Date: July 2008
Software Release Date: July 2008

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

Trademark Notices

All marks mentioned in this document are the property of their respective owners.

Updating documentation

Documentation enhancements are a continual project at Hewlett-Packard Software. You can update the documentation set at any time using the following procedure (which is also available in the HP OO readme file).

To obtain HP OO documentation

1. On The Web site https://support1.opsware.com/support/index.php, log in with account name and password that you received when you purchased HP OO.
2. On the **Support** tab, click the **Product Docs** subtab.
3. Under **Quick Jump**, click **Operations Orchestration** (or **Process Automation System**).
4. Under **Operations Orchestration**, click **ZIP** beside **HP OO 7.20 Full Documentation Set**.
5. Extract the files in the .zip file to the appropriate locations on your system:
   - For the tutorials to run, you must store the .swf file and the .html file in the same directory.
   - To obtain the repository that reflects the state of the flow at the start of the tutorial, unzip the file Exportof<preceding_tutorial_name>.zip.
   - To obtain the scriptlet for the tutorial that includes using scriptlets, click the scriptlet .txt file name.
   - To update your Central or Studio Help:
     a. Under **Help Files**, and then click **Studio Help File Bundle** or **Central Help File Bundle**.
     b. In the **File Download** box appears, click either **Open** or **Save**.
     c. Extract the files to the Hewlett-Packard Software\HP OO home directory, in either the **\Central\docs\help\Central** or **\Studio\docs\help\Studio** subdirectory, overwriting the existing file.

Where to Find Help, Tutorials, and More

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 runs

  The 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 flows

  The 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.

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 who 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.

Support

For support information, including patches, troubleshooting aids, support contract management, product manuals and more, visit one of the two following sites:

- https://support1.opsware.com/support/index.php
- http://www.hp.com/go/hpsoftware/DCA_support

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 HP 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, 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 the 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 (http://<your-hostname>) to the Intranet Security Zone, using the default settings.

To start Central

1. Start your machine's Web browser.
2. In the Web browser, add the Central site to the Trusted Sites security zone and to make sure that the Trusted Sites zone is set to its default level of security.
3. Paste the URL that your administrator sent you into the **Address** box of the Web browser and then press Enter.
4. When the message appears that you are about to view pages over a secure connection, click **OK.**

    A message appears, warning you that the site is not trusted. However, it is safe to proceed.
5. Click **Yes.**
6. 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.

You can customize the Central Dashboard to fit your needs, whether you run flows, administer HP OO, or manage IT.

After you log out or shut down the system, the settings revert to the default. The default setting for **flow Metrics** is to show the number of runs over the last seven days. **Popular flows** orders the flows listed there by number of runs.

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.
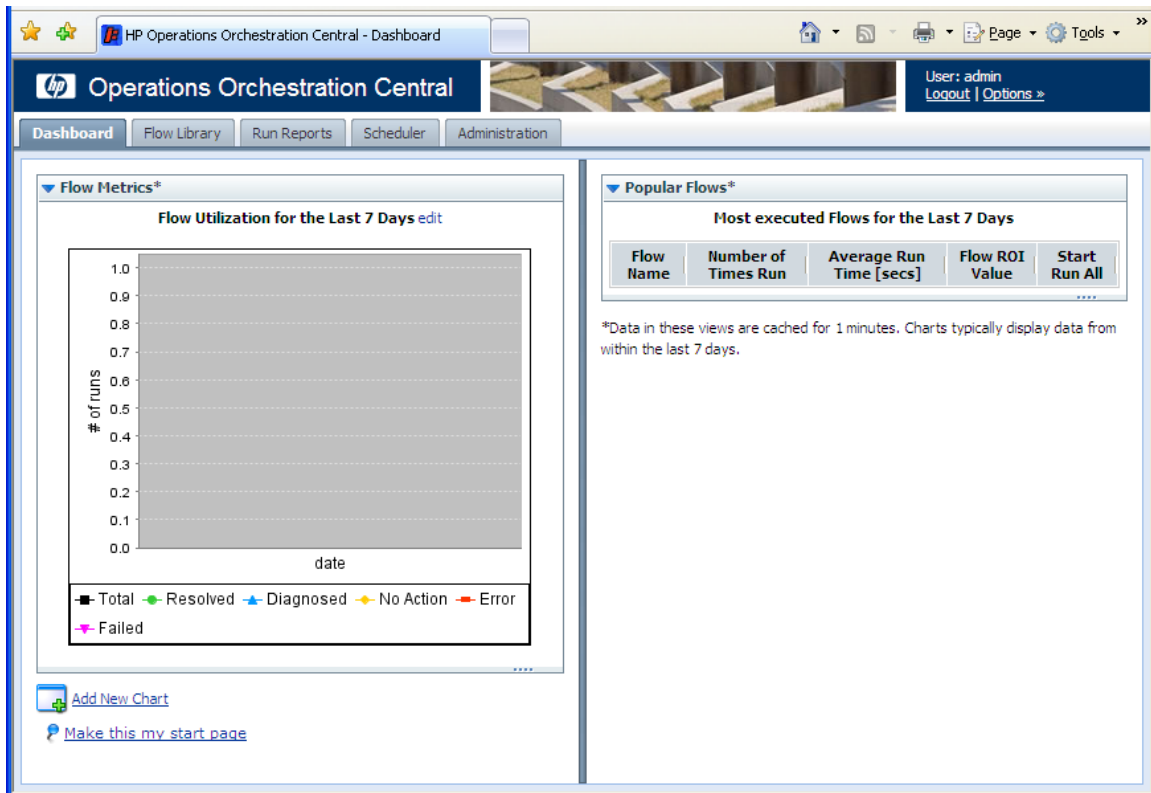
Figure 1 - Central Web application Dashboard

The **Flow Metrics** area is a diagnostic and analytical 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.

Central provides a graphical user interface for:

- Finding and running flows.
- Creating reports and viewing information on 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 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.
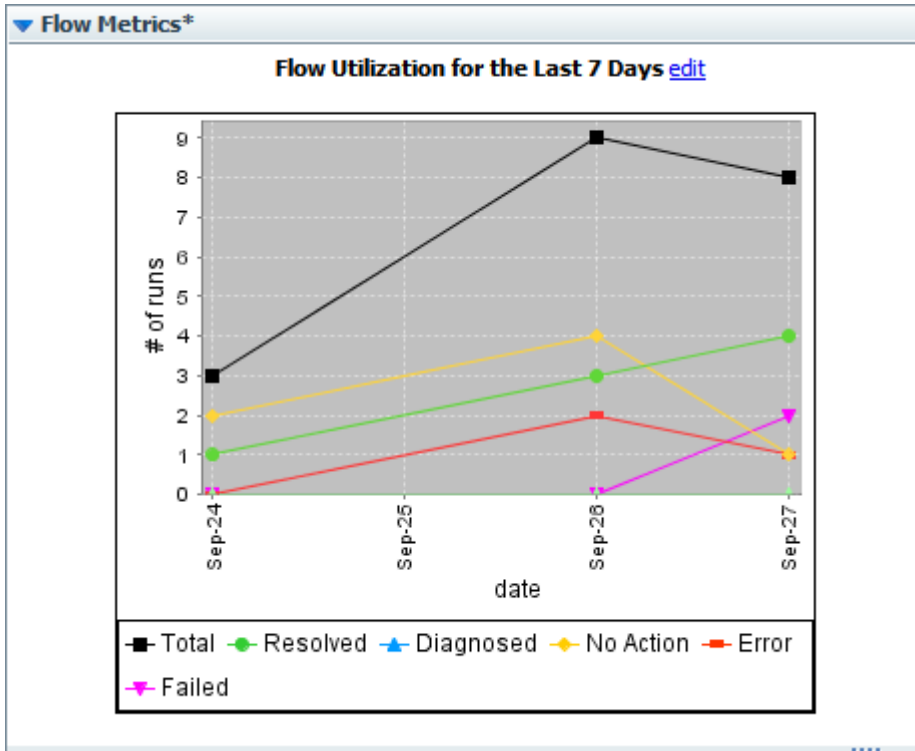
Figure 18 - Flow Metrics area

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.

Figure 19 - 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

Central 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, Run Reports, Scheduler,** and **Administration** tabs.
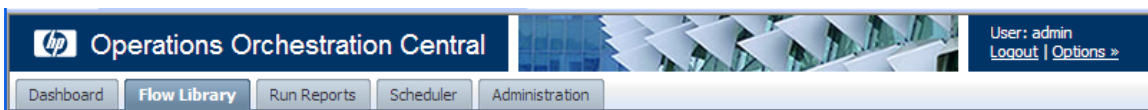
Figure 20 - Navigation tabs

Changing your HP OO Central password

To change the password for the current user

1. At the right end of the Central banner, click **Options**.
2. In the context menu that appears, click **Change Password**.
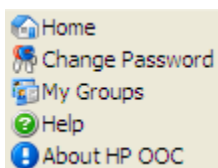
Figure 5 - Options menu

The **Change Password** dialog box appears.

3. In the **Current Password** text box, type your current password.
4. In the **New Password** and **Confirm New Password** text boxes, type the new password.

Viewing the HP OO groups that you belong to

To view the group memberships of the current user

1. At the right end of the Central banner, click **Options**.
2. In the context menu that appears, to view the groups that you are (that is, that the logged-in user is) a member of, click **My Groups**.
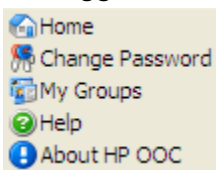
Figure 6 - Options menu

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.
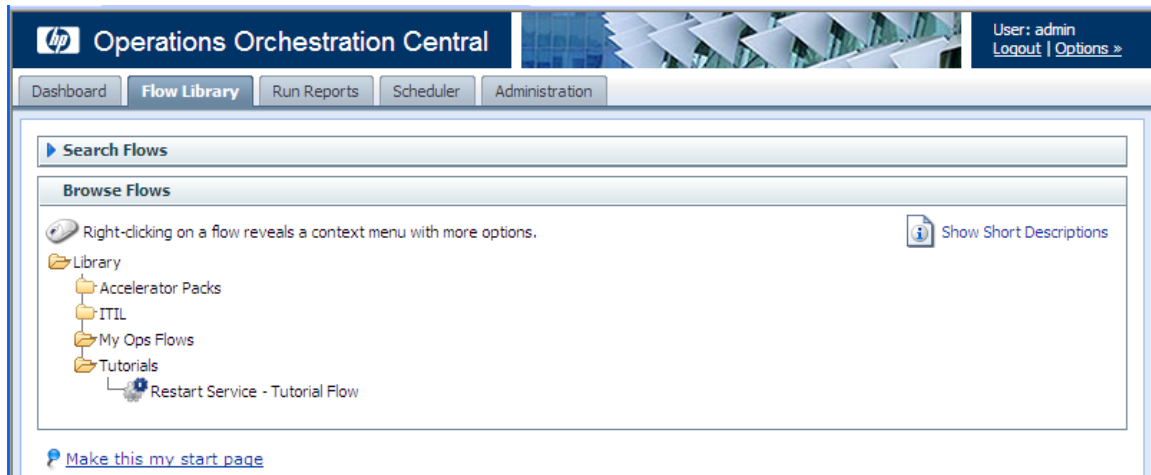
Figure 2 - Library

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 IIS SMTP server health, you would expand the **Accelerator Packs** folder and **IIS** folder, then run one of the **Server Health** flows.

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.**)

   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. The data remains only as long as the box is open.

   For more on running a flow, see <u>Three ways to run a flow</u>.

3. For more information on the flow, click the "i" balloon (ⓘ).

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

4. 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.

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** 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 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 universally unique ID (UUID) and 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
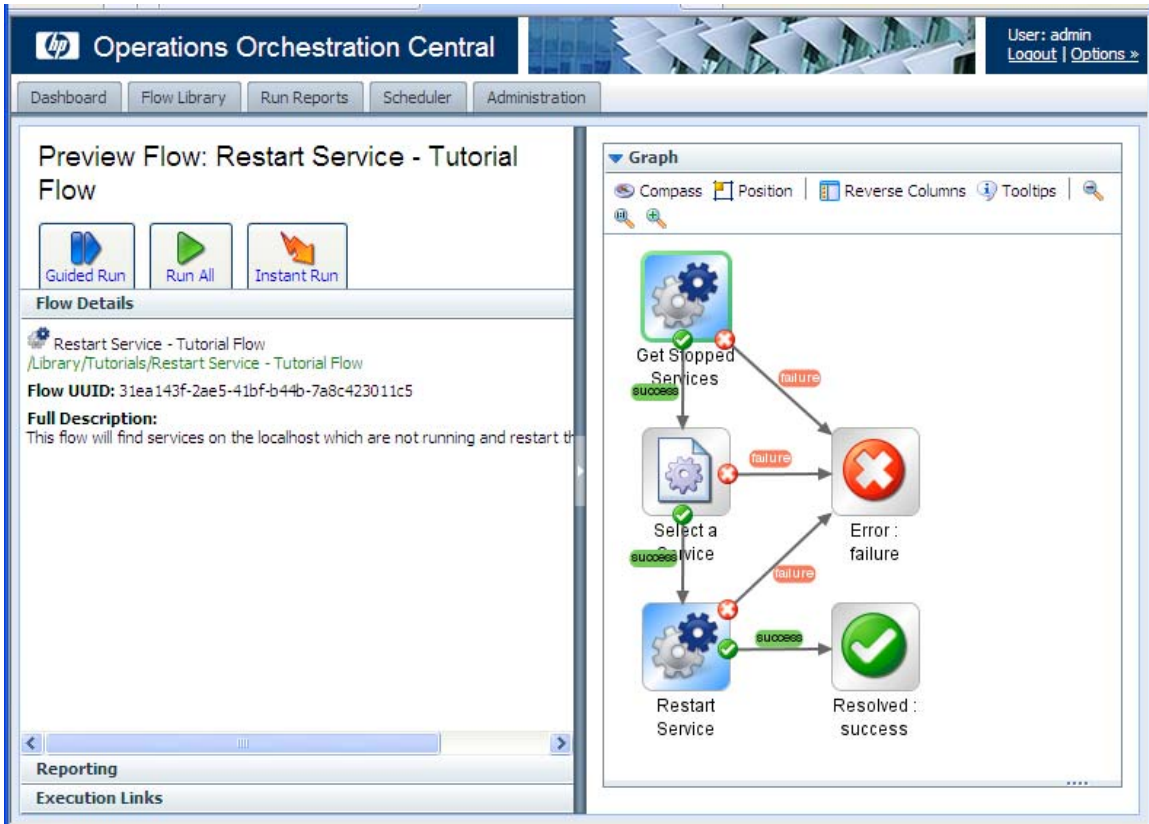
Figure 4 - The Restart Service flow, loaded and ready to run

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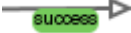
- To return to the Library without running the flow that you are previewing, just click the **Flow Library** tab again.
- To run the flow repeatedly, you can create a schedule for it. For information on scheduling repeated runs of a flow, see Scheduling flows.

The following table describes the symbols that may appear in 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.
	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

	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
success	Gated transition	A transition that is gated, or restricted to users with certain access permissions, appears in red on the canvas.
success	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 9 - 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 ( ).

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, within which the new run starts and runs to completion.
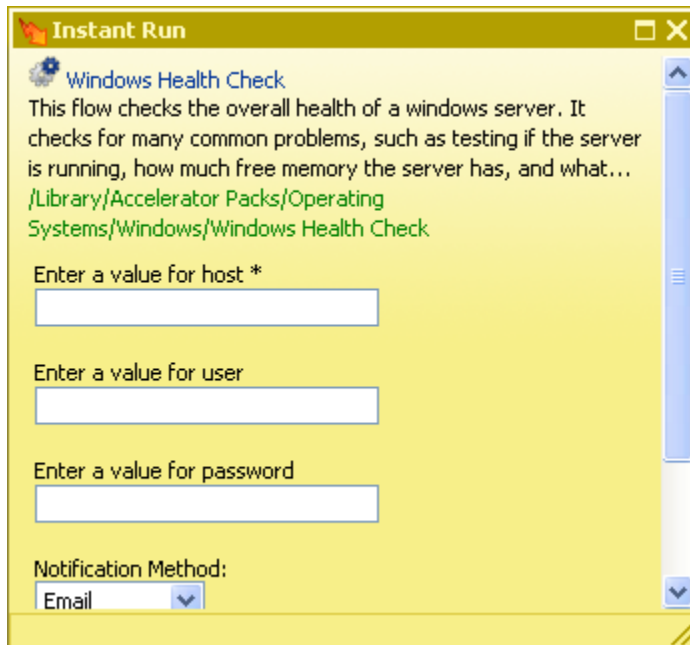
Figure 10 - Instant run of a flow

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.
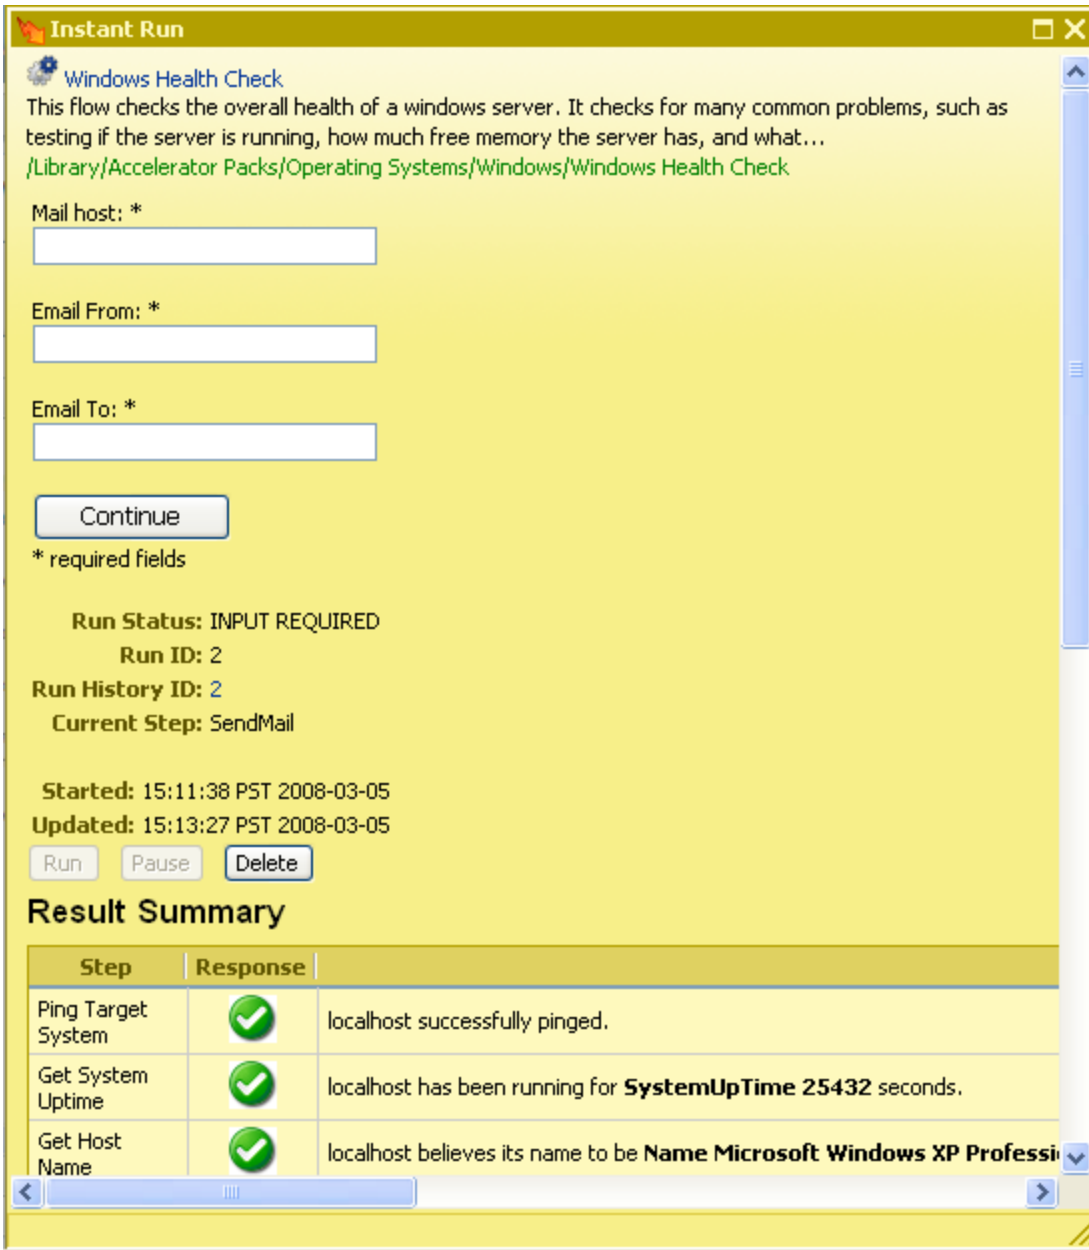
Figure 11 - Instant Run dialog box, expanded

The flow runs to completion on its own, unless a user prompt needs a response.

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 run you want to flow.
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**

Running subflows

When you are running a flow step by step 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.

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.
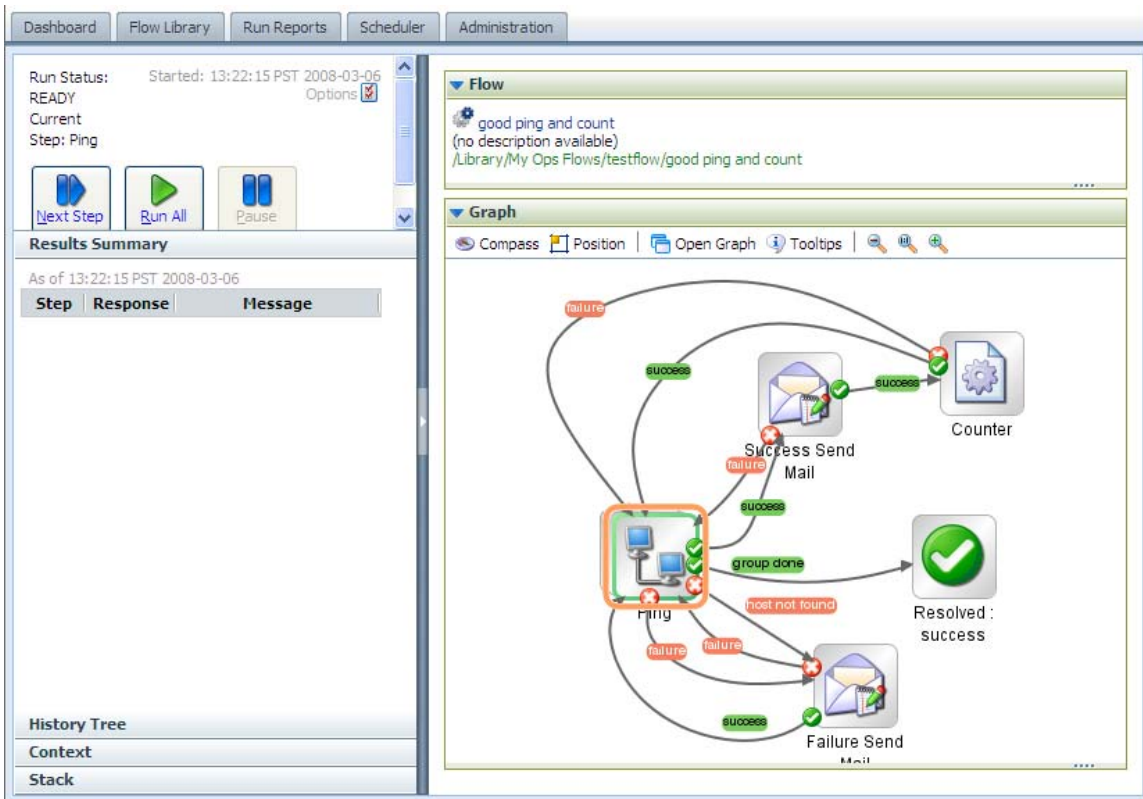
Figure 12 - Flow graph in same window

Here is the same flow after you click **Open Graph**.
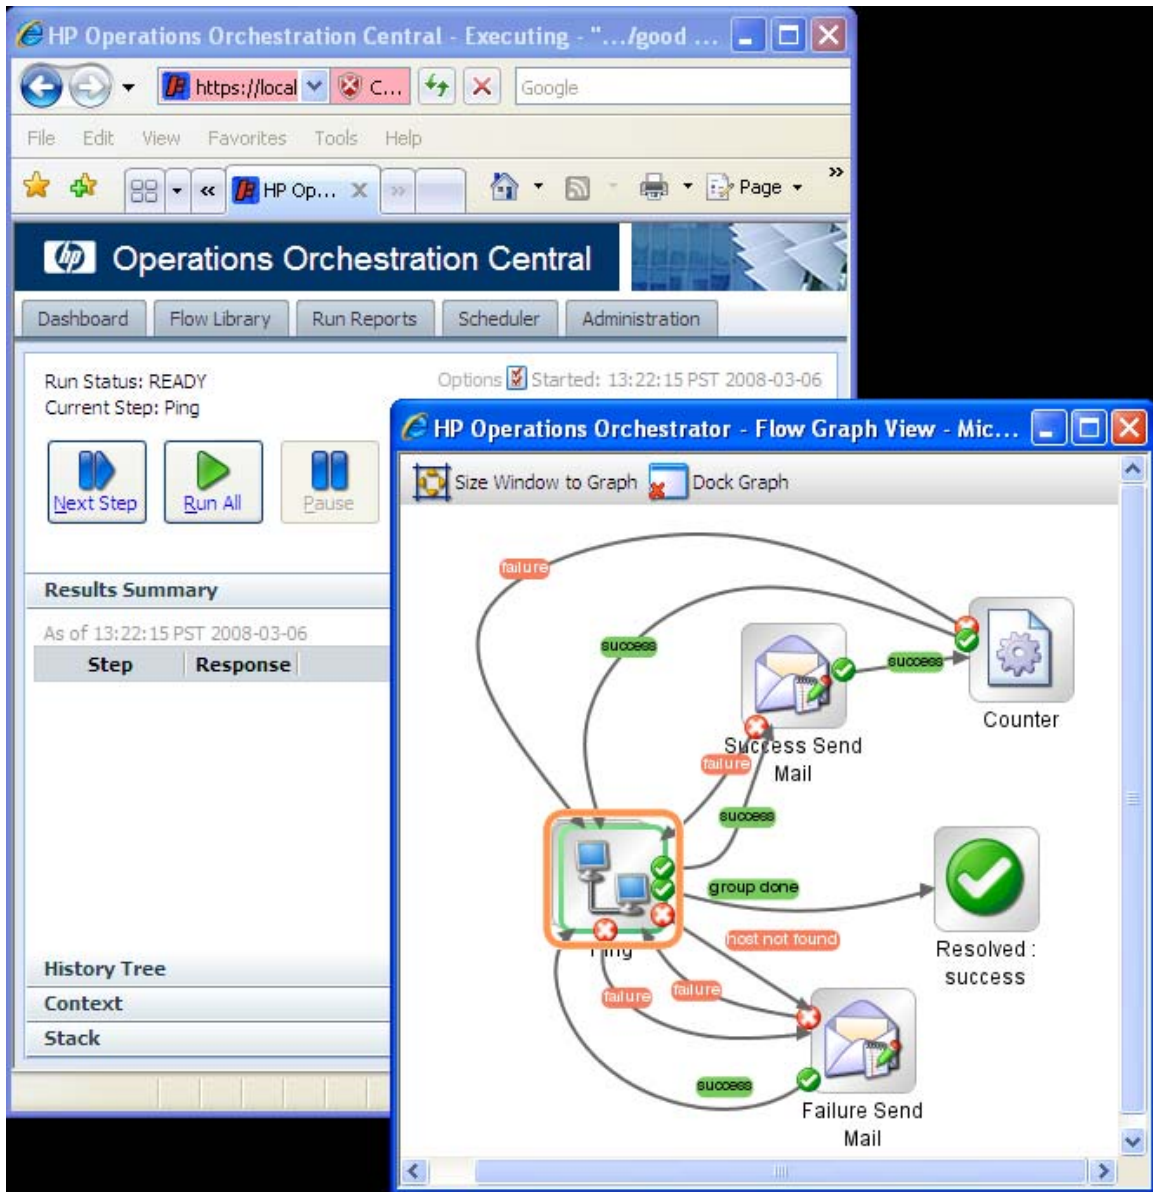
Figure 13 - Flow graph undocked from Library tab

Providing even a small flow with its own browser window provides much more room for the **Results Summary**.

To open the flow graph in a different browser window

1. With the flow running in step-by-step 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.

Figure 14 - Results Summary panel, mid-run

Following is a close-up of the Results Summary.

Figure 26 - Results Summary of a flow run

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

To view a flow's run history mid-run

1. On the **Administration** tab, under **Name**, click the downward-pointing arrow by the name of the flow whose run you want to delete.

   A drop-down menu appears, including the command **Inspect History**.
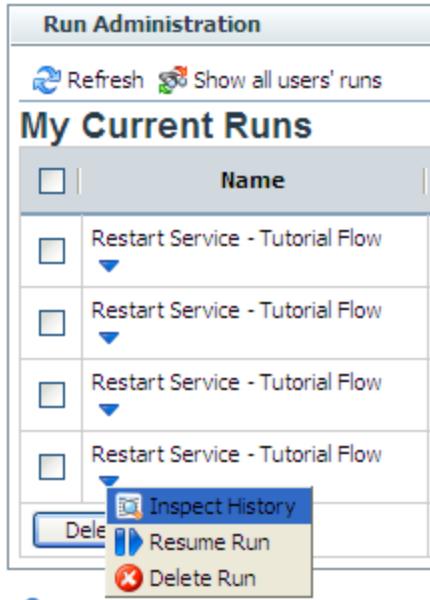
Figure 27 – Runs in the Run Administration panel

2. Click **Inspect History**.

   The run history is a report for the single run, which opens on the **Run Reports** tab.
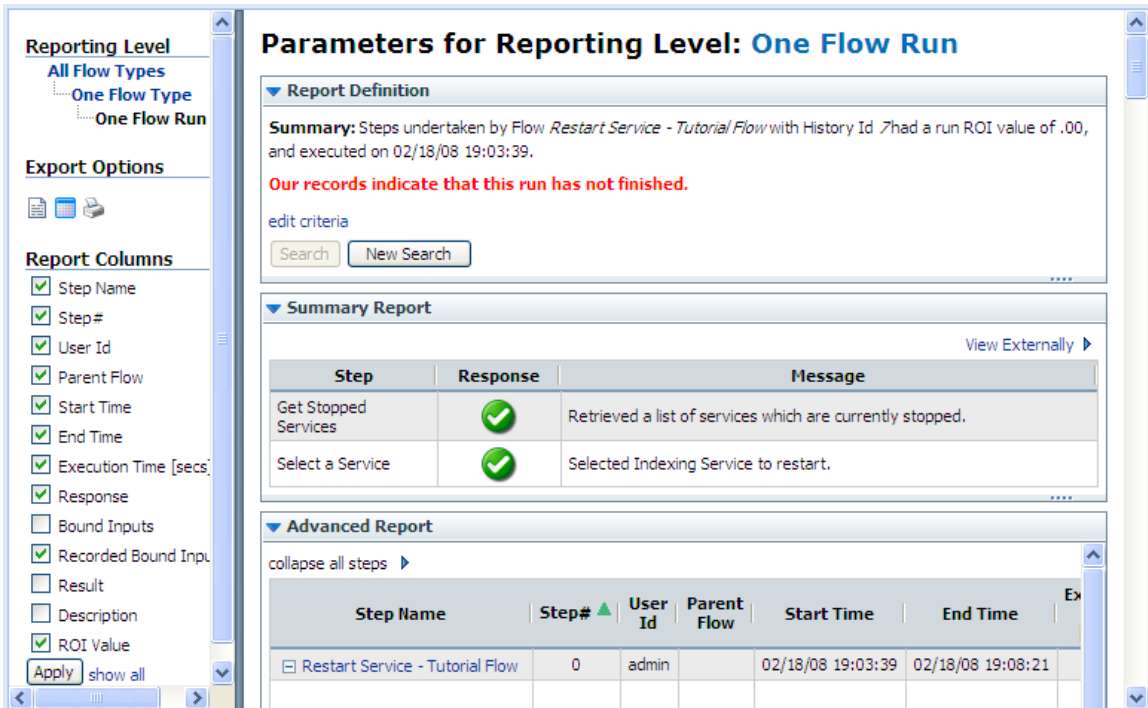
Figure 17 - Run History

The Summary Report that appeared where you run the flow also appears here in the Run Report. Information in more depth on what has happened in each step that has completed is under **Advanced Report**.

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

3. To select which columns appear in the **Advanced Report** section, select the column names under **Report Columns**, and then 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

Suppose that on the Dashboard charts you have identified server 192.138.16.133 as having crashed frequently and you want to look at what happened in the flows that were run against it. Now that you have zeroed in the server as a problem area, you can examine run histories from the highest level down to what happened on an individual step in an individual run.

On the Dashboard, by looking at the **Flows per Configuration Item** chart, you could identify the flow that was run to diagnose the server's problem and bring it back online. For example, suppose that this flow was the Restart Service flow.

To look in greater detail at what the Restart Service flow did when it restarted the server, on the **Run Reports** tab, you could define a search that finds all runs of the Restart Service flow, then drill down to the run that fixed the server. Within that run, you could then examine the **Summary Info** column for each run.

Parameters for Reporting Level: **All Flow Types**

Figure 14 - Defining reports of run histories

To define which run histories are examined

1. Specify the time window over which you want to see run histories.

   For a fixed window, pick start and end dates by clicking the calendars. Let's click the calendars and pick 09/04/06 for the **Start** date and 09/06/06 for the **End** date.

   If you specify the time, do so in 24-hour format.

2. Under **Filters**, in the **Subtree** drop-down list, pick the Library path that contains the flow whose runs you want to see.

   The drop-down list contains only the subtree paths that contain flows that have been run.

   In our example, we'll pick **/Library/Accelerator Packs/Windows Management**.

   Although for this example we'll omit specifying users, you could do so, separating the user names with commas.

3. To specify a flow run or modified by a user, type the user's account name in the Executed by User(s) or Last Modified by User(s) box.

   If you type more than one user in a box, separate the account names with a comma.

4. Select the **Result** check box or boxes that you want the report to include.

5. Under **Matching Step Inputs**, add any input names and values you want to use to further limit the runs included in the report.

6. To specify another step input name and value, you click **+1 More**, then repeat the preceding step, specifying the second server's IP address or name.

7. To use one or more domain terms to further define the search, then under **Matching Domain Values**, follow the same steps to specify domain term name/value pairs that you did for step inputs.

8. Click **Search**.

The search results appear.

▼ Report Data

Found 7 results at 06/01/07 11:24:12

Ops Flow Name	Number of Runs ▼	Avg. Number of Steps	Avg. Repair Time [secs]	Most Recent Run	Percent Resolved	Percent Diagnosed	Percent Error	Percent No Action	Percent Failed	Flow Value
Patch Process	5	6.00	.196	05/31/07 13:12:54	0.00%	60.00%	0.00%	40.00%	0.00%	$83.83
Change Notice	4	6.00	.343	05/31/07 13:09:59	25.00%	50.00%	0.00%	25.00%	0.00%	$0.00
Infrastructure Alert	3	7.00	.434	05/31/07 13:11:14	33.33%	66.67%	0.00%	0.00%	0.00%	$53.47
SQL Manager Statistics	2	7.00	34.873	05/31/07 12:21:18	100.00%	0.00%	0.00%	0.00%	0.00%	$0.00
Application Alert	1	7.00	.862	05/31/07 13:07:24	100.00%	0.00%	0.00%	0.00%	0.00%	$16.17
Restart Service - Tutorial Flow	1	90.00	10.285	05/31/07 11:45:15	100.00%	0.00%	0.00%	0.00%	0.00%	$2.15
Simple SMTP Check	1	4.00	62.334	05/31/07 11:48:47	0.00%	0.00%	100.00%	0.00%	0.00%	$0.00

Figure 15 - Search results

9. To select the columns that are displayed, in the **Report Columns** panel on the left of the search results, unselect the columns you're not interested in.

Figure 20 - Selecting the run-history data types to display

Run histories: single flow, single run

You can narrow the run history to a single flow and a single run.

To look at run histories for a single flow and a single run

1. To see runs for a single flow, click the flow's name.

▼ Report Data

Found 5 results at 06/01/07 11:35:08

Flow Id	Start Time ▼	User Id	Time To Run [secs]	Run Result	Number of Steps	Flow Path	Run Value
50	05/31/07 13:12:54	admin	.220	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	$16.43
49	05/31/07 13:12:37	admin	.211	No Action Taken	6	/Library/My Ops Flows/Dashboard/Patch Process	$17.27
48	05/31/07 13:12:19	admin	.160	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	$16.43
47	05/31/07 13:12:03	admin	.181	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	$16.43
46	05/31/07 13:11:41	admin	.211	No Action Taken	6	/Library/My Ops Flows/Dashboard/Patch Process	$17.27

Figure 16 - Run history of one flow

2. To see a single run, click the run number.

The data are the **Summary** and **Advanced** reports.

Figure 17 - Run report for a single flow run

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

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.

Flows that you create schedules for 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 server whose online status Server Status Flow should check. 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.

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 Administration Guide* (AdminGuide.pdf).

Creating a schedule for a flow

To create a schedule for a flow

1. Right-click the flow on the **Flow Library** tab…
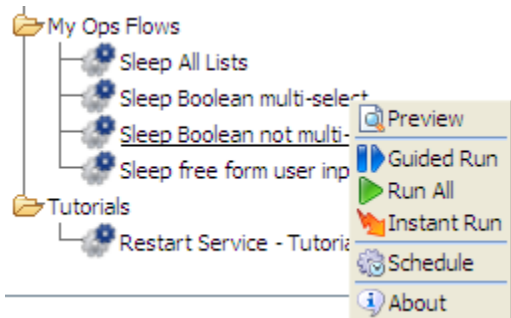
   …and click **Schedule**.

2. In the following box, on the **Schedule** tab, specify the time(s) that you want the flow to run in the boxes.

3. On the **Inputs** tab…
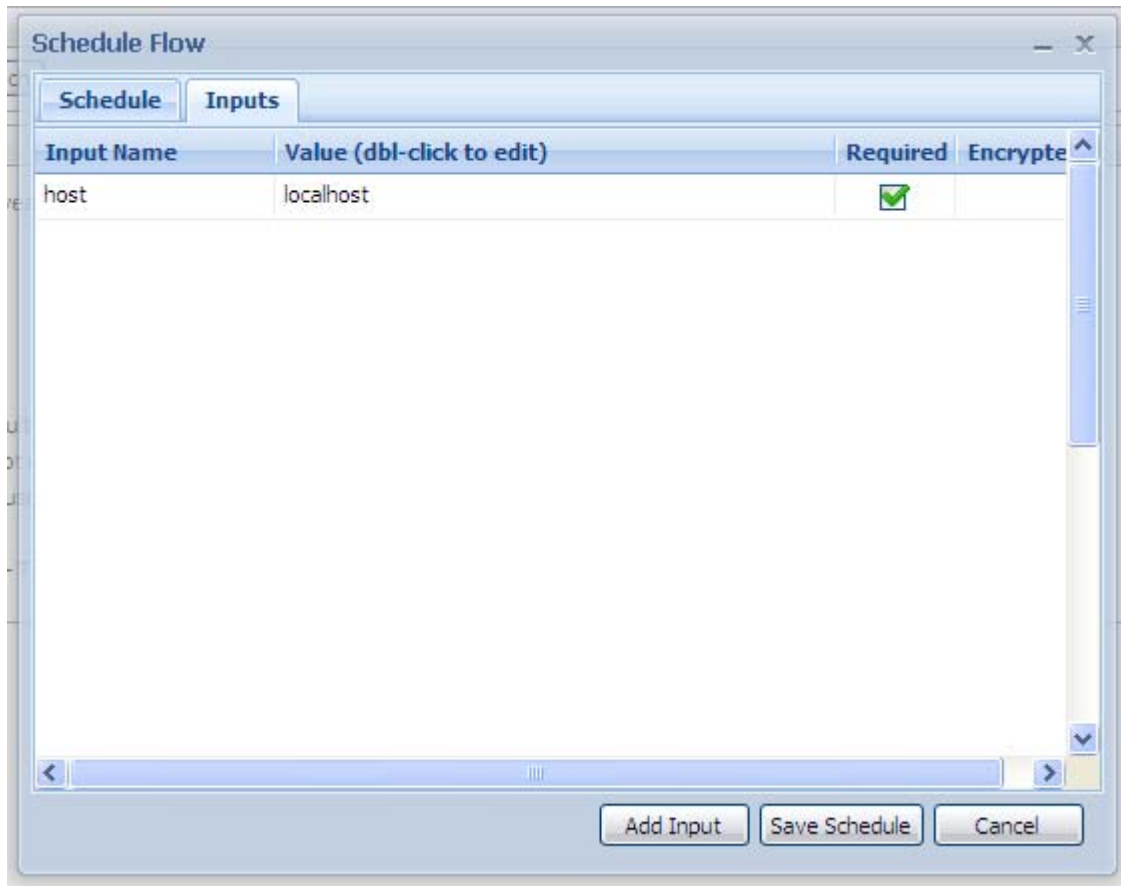
…fill in the **Value** for all of the specified flow input names.

The flow's inputs are automatically listed in the **Name** boxes.

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 were the sequence of a comma followed by a space.

4. To add an input:

    - On either 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.

5. On either tab, click **Save Schedule**.

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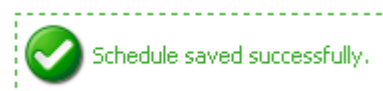
Once you have created a schedule, you can edit it on the **Scheduler** tab.

Viewing and editing flow schedules

The main page of the **Scheduler** tab lists the flows that have schedules created for them.
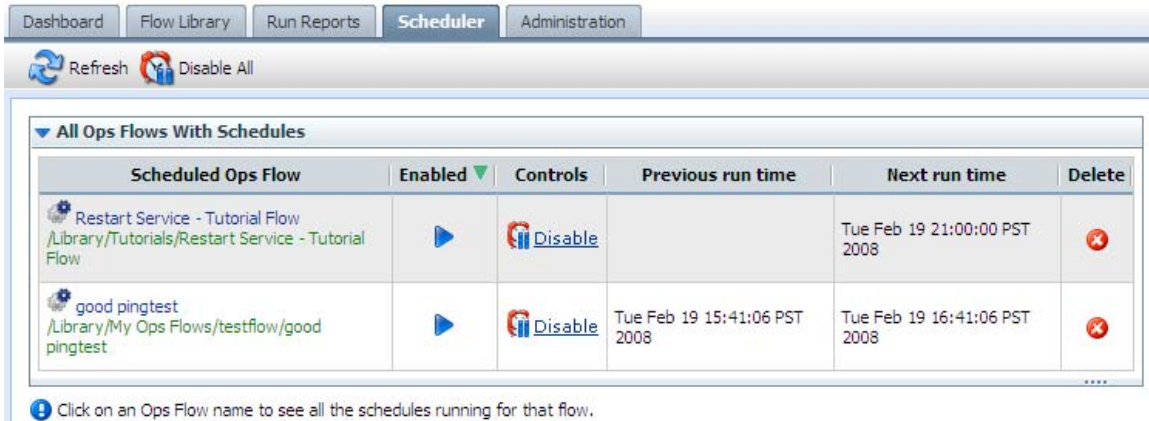
Figure 5 — Scheduled flows on the Scheduler tab

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.
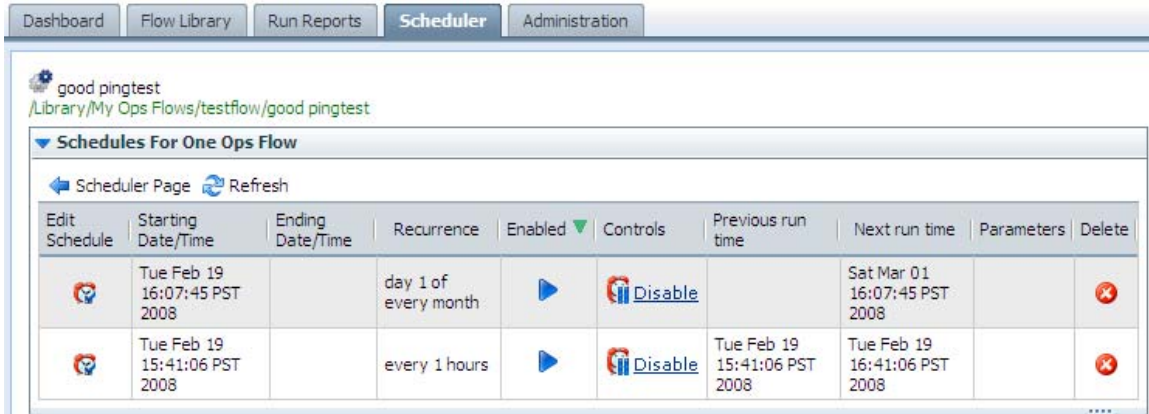
Figure 6 - Schedules for a single flow

Note that you can edit the schedule or inputs by clicking the clock icon in the **Edit Schedule** column.

Editing existing schedules

To edit a flow's schedule

1. To edit a schedule, on the **Scheduler** tab, select the schedule, and then click the 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

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**, as appropriate.
2. To enable or disable a single schedule, on the **Scheduler** tab, select the schedule, and then click **Enable** or **Disable**, 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 **Administration** tab. To see the **Administration** tab, you must be logged in to Central with an account is a member of the ADMINISTRATOR group.

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:

Figure 7 – 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		Because Scheduler is entirely part of Central,

"checks in" with the other instances of the cluster…		this is redundant with the clustering configuration. Take no action.
If Scheduler(s) are clustered		Because 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 **Run 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.
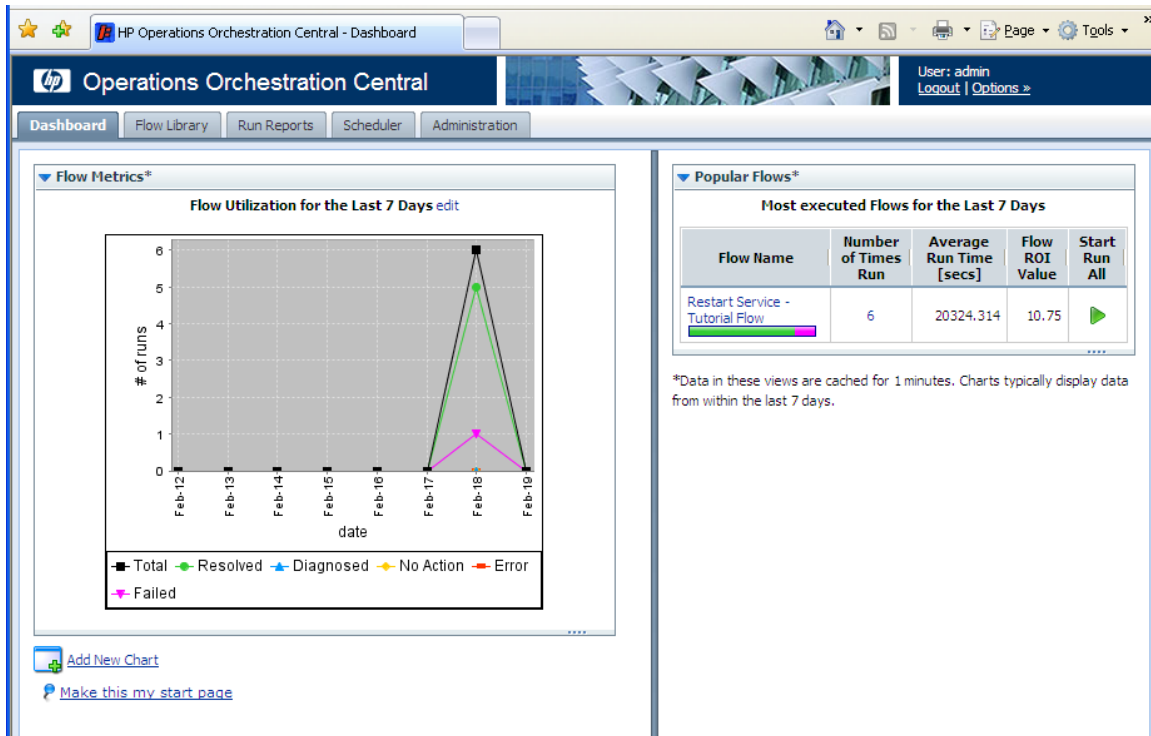
Figure 26 - The Dashboard

Charts that you can bring up on the **Dashboard** tab can tell you:

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

    Note: HP OO uses the ITIL term Configuration Item (or CI) to refer to server, 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.

New Chart

Select a report to view

Alerts per Configuration Item for the Last 7 Days
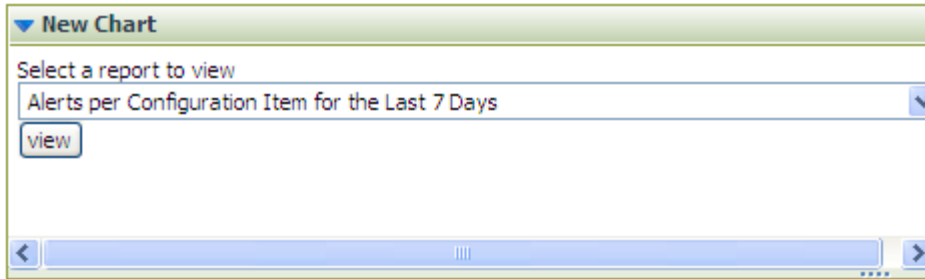
view

Figure 27 - 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

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, No Action Taken, Resolved, and Failure).
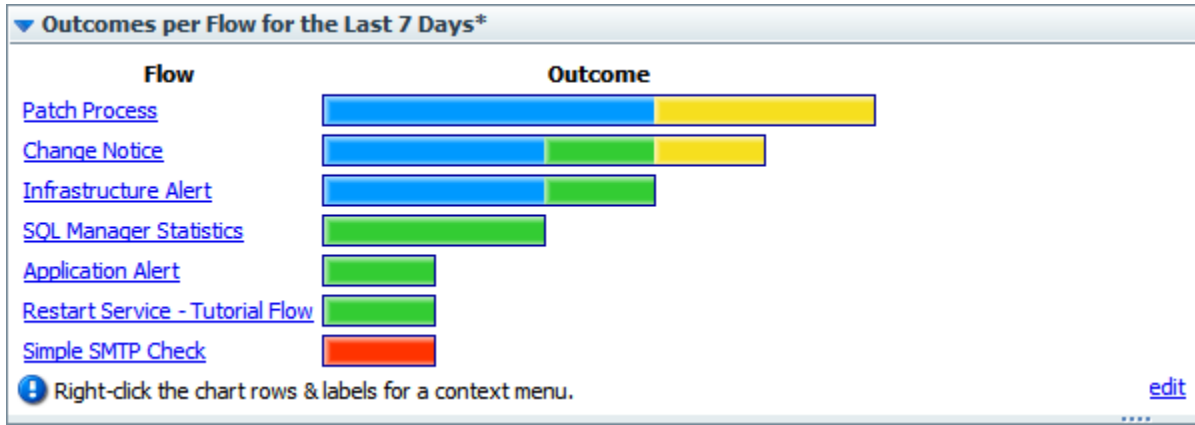
Figure 8 - Sample Dashboard chart

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.
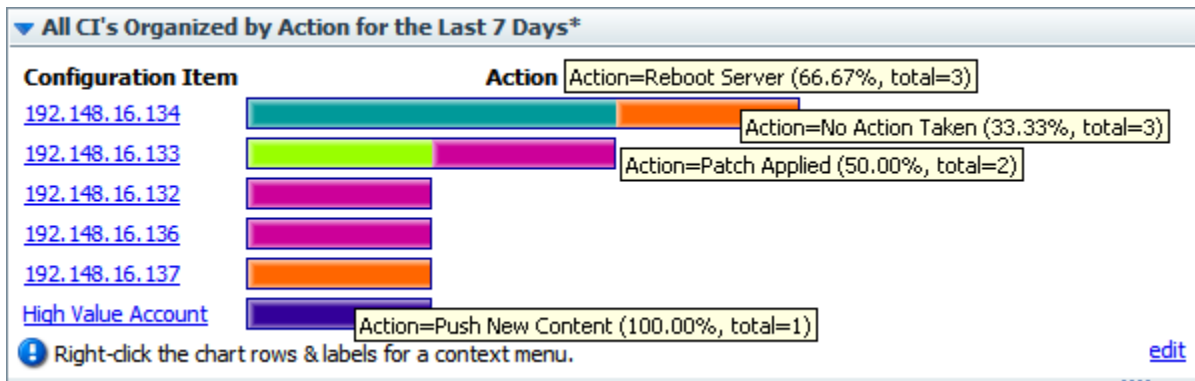
Figure 9 - Sample Dashboard chart with composite of bar labels

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 192.148.14.152 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.

Configuration Item   **Action**

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.
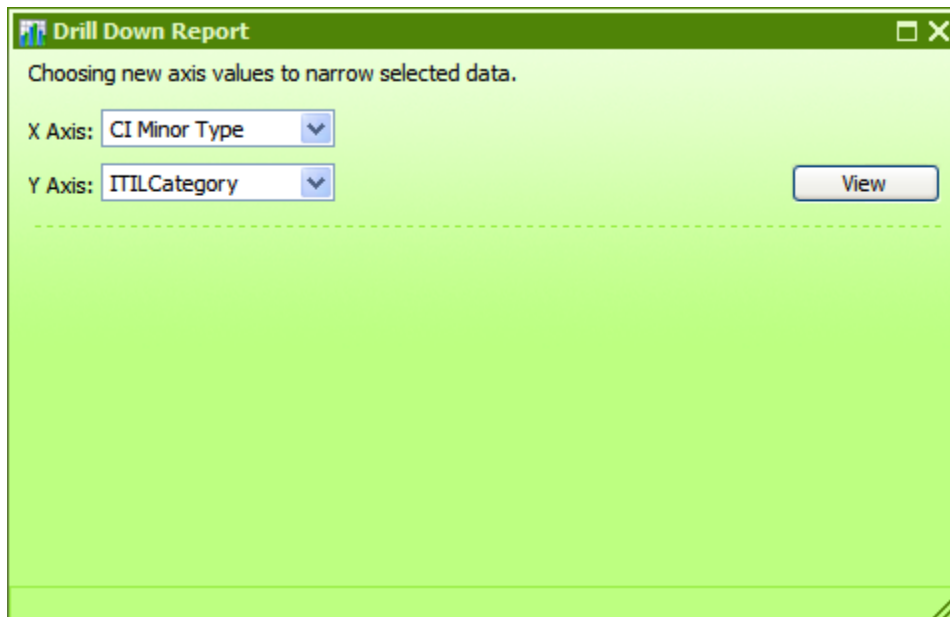
   Figure 10 - Creating a drill-down report

2. 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.
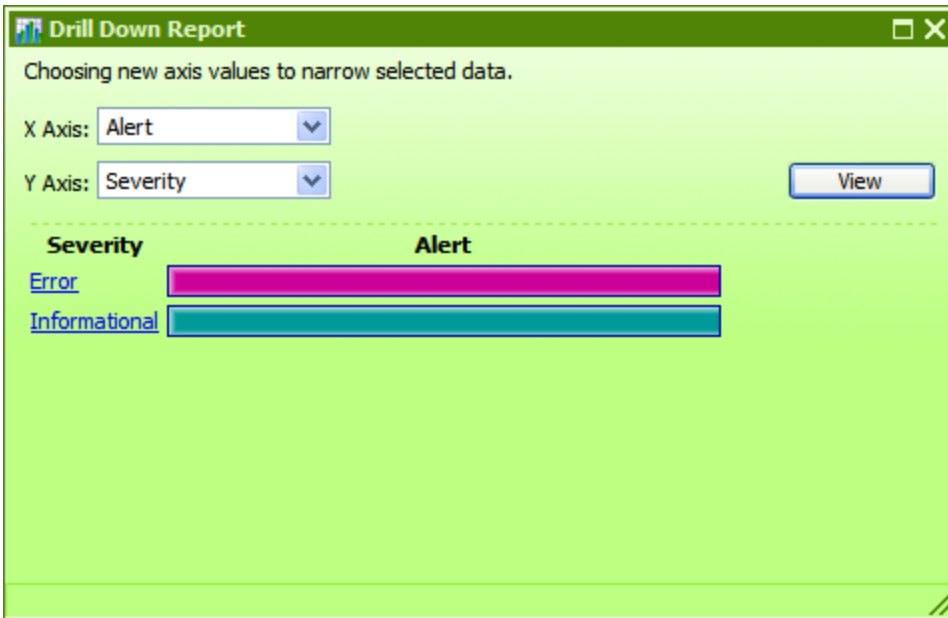
Figure 11 - The drill-down report that we created

3. 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.

4. 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 there are 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.

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.

Figure 12 - Administration tab, Dashboard chart definitions

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

   A new chart definition box appears.

Figure 13 - 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 34 - Chart definition, Advanced Details

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

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. Under **Additional Restraints**, 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.

12. 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, 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 Domain Terms section in the Help for Studio.

Exporting and importing chart definitions

You might want to copy chart definitions from one installation of Central to another when:

- You have similar IT health and status concerns across multiple organizations or domains and so need consistent reporting and metrics in each area.
- You have clustered Central for failover and want continuity with the Dashboard charts when a Central server goes down and another in the cluster picks up its load.

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.

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.

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 several ways to start a flow from outside the Central application. Of the following three ways, using Rsflowinvoke, is the recommended method.

- Using Rsflowinvoke, a command-line tool that comes in a Windows and a Java version

  For more information, see <u>Starting a flow using Rsflowinvoke</u>.

- Building a URL that you then paste into a browser

  For more information, see <u>Building a URL for starting a flow</u> and <u>Starting a flow from a URL</u>.

- On a Linux system, using wget

You can also search for a flow from outside the Central application. For more information, see .

Both Rsflowinvoke and wget use the URL that you build. To build the URL, you must define a prefix for each input (parameter) used in the flow.

Defining a prefix for inputs in URLs that launch flows

In URLs that start flow runs, initial inputs must not be named "service" or "sp". To protect your flow runs from errors that can result from use of these reserved names, you can define a prefix that is required for all initial-input names used in the URL. If you do not define such a prefix, the initial-input names do not need to have any prefix.

Important: Do not specify as your prefix a character that is reserved or could be misunderstood when used in URLs.

To define a prefix for initial-input names in URLs that start flow runs

1. Log in to Central with an account that has HP OO administrative rights.
2. On the **Administration** tab, click the **System Configuration** subtab.
3. In the **General Settings** area, in the **Prefix for init params** row, in the **Value** box, type a prefix that you will use in the URL that you send to your recipient.
4. Click **Save General Settings**.
5. Restart the HP OO Central service (RSCentral).

Creating an encrypted password

To create an encrypted password for use with Rsflowinvoke

- Type and execute the following command:

  `Rsflowinvoke.exe –cp`

  You are prompted to type and then repeat the password. The password that you type is then encrypted. When you then run Rsflowinvoke.exe with the encrypted password that you have just created, you use the `–ep` option with the encrypted form of the password.

Starting a flow using Rsflowinvoke

RSFlowInvoke (Rsflowinvoke.exe or the Java version, Jrsflowinvoke.jar) is a command-line utility that substitutes for a Web-browser call in starting a flow. Substituting a command-line for a Web-browser call enables you to start a flow without using Central (although the Central service must be running). RSFlowInvoke (or JRSFlowInvoke) is useful when you want to start a flow from an external system, such as a monitoring application, that can use a command line to start a flow.

Rsflowinvoke.exe can take a URL as an argument, or you can supply the host's identity, port, and flow's identity with the Rsflowinvoke command's options.

You can run Rsflowinvoke from any machine from which you can log into Central (https to port 8443). This is why you can run Rsflowinvoke from a monitoring program such as MOM.

You can use Rsflowinvoke:

- From a command-line window
- As part of a script or batch file
- From any application that can use a command line

Both the Windows and Java versions of Rsflowinvoke are available in the HP OO home directory, in the \Central\Tools folder.

To start a flow using Rsflowinvoke in a batch file

- In the script or batch file, use one of the following syntaxes for starting a flow with Rsflowinvoke:

  - If you are using RSFlowInvoke.exe:

    `RSFlowInvoke.exe –host -flow -inputs –u  -p -a -ep -rc -rw -t -v -verbose]%1`

  - If you are using Jrsflowinvoke.jar

    `-jar JRSFlowInvoke.jar –host -flow -inputs -u -p -a -ep -rc -rw`

    Where:

    `-host`

    is followed by the hostname and port number, separated by a colon (`hostname:portnumber`).

    `-flow`

    is followed by the flow name or UUID.

    `-inputs`

    is followed by the inputs, using the pattern `name=value&name2=value2`

    `-u`

    is followed by the username.

    `-p`

    is followed by the password. If you use –ep to specify an encrypted password, do not use this option.

    `-ep`

    is followed by the encrypted password. If you use –p to specify a non-encrypted password, do not use this option.

`-a`

is followed by the authentication type (`Basic` or `Digest`).

`-t`

is followed by the timeout, in seconds. The default value is 100 seconds.

`-v || -verbose`

means that all output is written to the screen.

`-rc`

is followed by the number of times to retry a flow that fails for any reason. The default is 0; the maximum is 30.

`-rw`

is followed by the number of seconds to wait between retries. The default is 5 seconds.

To reference the flow by a URL, replace the –host and –flow values by the URL for the flow. For information on building the URL, see Building a URL for starting a flow.

For creating an encrypted password from within Rsflowinvoke, see Creating an encrypted password.

Finding out what happened in Rsflowinvoke

The following return codes tell you what happened in the running of the flow from Rsflowinvoke:

- 0 - The flow was run. This code is not related to the flow's response.
- 1 - Central responded with HTTP code 503. This normally means it lacks the resources needed to run the flow.
- 2 - An Unknown internal server error occurred in Central.
- 3 - RSFlowInvoke was unable to authenticate against Central.
- 4 - The specified flow was not found, or the supplied URL was not found.
- 5 - A Socket Timeout occurred.
- 6 - An unknown socket (Communication) error occurred.
- 7 - An unknown error occurred.

Registering Rsflowinvoke with the Global Assembly Cache

The Global Assembly Cache (GAC) is a store on a local .NET machine for assemblies of .NET code. If you register Rsflowinvoke.exe with GAC, you can start the flow from within a .NET application, using any .NET-compatible language, such as C#.

To register or unregister in GAC

1. On a .NET machine, open a command window and type a command with the following syntax:

   `gacutil.exe [/i|/u] RSFlowInvoke.exe`

   Where:

- /i registers Rsflowinvoke.exe with GAC.
- /u unregisters Rsflowinvoke.exe with GAC.

2. Once Rsflowinvoke.exe is registered with GAC, type the following to view the assembly information:

   `RSFlowInvoke.exe -s`

Building a URL for starting a flow

A URL of the correct form is necessary for starting a flow in a command-line utility or in the address box of a web browser.

The main variations in building the URL are:

- Whether you identify the flow by its name or by its universally unique ID (UUID).
- The initial inputs (flow input values) that are required for the flow to run. As when you run a flow automatically, starting a flow from outside Central requires that all the inputs are known before the flow starts, so no user intervention is required during the flow's run.

  Specifying initial inputs requires a prefix for identifying each name/value pair as an input. Do not use characters that are reserved when used in URLs or characters that can be misunderstood in URLs. For information on defining a prefix, see <u>Defining a prefix for inputs in URLs that launch flows</u>.

  If the flow you are starting has a multi-instance step, use a comma to separate the values for the input that has multiple values.

  Thus, if you define the sequence _xx_ as the prefix for initial inputs, a set of inputs might look like the following:

  `_xx_input1=localhost_xx_input2=8443_xx_input3=`

  If the flow has a multi-instance step, then for the step's input that has multiple values, you separate the multiple values with the separator character that the flow's author specified in Studio. Suppose that in the above set of input definitions input1 has two values (the IP addresses 10.0.0.100 and 10.0.0.101) and its separator character is a comma (,). The input definitions would look like this:

  `_xx_input1=10.0.0.100,10.0.0.101_xx_input2=8443_xx_input3=`

Note: You only have to specify values for inputs that:

- Get their value from user prompts.
- Have not been assigned a value (or a way to get a value).

  This means you do not have to specify a value for any inputs that have a specific value (or set of values, as in multi-instance steps or steps that get their value from an Iterator operation) assigned to them or for inputs that get their value from a system account or from the logged-in user's credentials.

The URL syntaxes for the two methods of identifying the flow are:

- Identifying the flow by folder path and name:

  `https://<Central host>:<Central port>/PAS/services/http/execute/Library/<folder path>/<flow name><?param1=value1&param2=value2...> -u <user> -p <password>`

- Identifying the flow by its UUID:

```
https://<Central host>:<Central
port>/PAS/services/http/execute/<flow
UUID><?param1=value1&param2=value2...> -u <user> -p <password>
```

Note: If the password is encrypted, the password switch is –ep instead of -p. For information on creating an encrypted password within Rsflowinvoke, see Creating an encrypted password.

You can also obtaining a return from a flow run without waiting for it to complete. This is called running the flow *asynchronously.*

To run the flow asynchronously

* In the URL, replace

  `/execute/`

  with

  `/execute_async/`

For an example, let's suppose the following:

* You're starting a copy of the Connectivity Test flow that you've modified so that the **Remote Ping** step is a multi-instance step with multiple values for the target input.
* The username and password inputs are assigned to a system account and the packets and packetSize inputs have specific values, so you only need to assign values to the host and target initial inputs.
  * The **host** input is localhost
  * The target input's values are 55.55.0.47 and 55.55.0.49
* The credentials for starting the flow are username "admin" and an encrypted password.

Here are two examples of URLs for launching this flow, using each method of identifying the flow, and an example of obtaining a return from a flow run without waiting for it to complete.

* Identifying the flow by folder path and name, the URL would be:

  ```
  https://81.123.17.68:8443/PAS/services/http/execute/Library/My
  Flows/Connectivity Test?&host=localhost&target=55.55.0.47,55.55.0.49
  ```

* Identifying the flow by its UUID, the URL would be:

  ```
  https://81.123.17.68:8443/PAS/services/http/execute/Library/
  e4ae39e0-30b5-4ef0-a08d-
  32f75dfa904f?&host=localhost&target=55.55.0.47,55.55.0.49
  ```

* To make the flow in the second example run asynchronously, the URL would be:

  ```
  https://81.123.17.68:8443/PAS/services/http/execute_async/Library/
  e4ae39e0-30b5-4ef0-a08d-
  32f75dfa904f?&host=localhost&target=55.55.0.47,55.55.0.49
  ```

Starting a flow from a URL

Besides manually building a URL, you can create the URL from within Central.

To create a linked URL that can start a flow run

1. Click the **Flow Library** tab, navigate to the flow, and click the flow name to open the preview of the flow.
2. Under **Execution Links**, select the URL in the text box of the desired type of run (either **Guided Run** or **Run All**).
3. To send the URL to another Central user, paste the URL into a message.

   OR

   To start the flow from outside Central, use the URL. For more information, see [Starting a flow from outside Central](#).
4. If the flow has any required inputs, modify the URL by adding name-value pairs that define values for all the inputs.

   For more information:

   - On defining a prefix for input-value pairs, see [Defining a prefix for inputs in URLs that launch flows](#).
   - On specifying values for initial inputs and when you need to do so, see [Building a URL for starting a flow](#).

Searching for a flow from outside Central

You can use the following syntax to search for a flow from outside the Central application. You must include a properly formatted query string that has a permissible search string.

```
java -jar JRSFlowInvoke.jar
https://<host>:<port>/PAS/services/http/search? queryString =
<sequence_of_term_expressions> -u <user> -p <password>
```

where:

- `<host>` is the Central server on which the search is performed.
- `<port>` is the port number on the Central server that Central uses to communicate.
- `<sequence_of_term_expressions>` is a single `term`

  OR

- a sequence: `term + (operator + sequence_of_term_expressions)`
    - `term` is `[field:]non_field_term`
        - The optional `field` is one of the fields described in [Searching for a flow](#).
        - `non_field_term` is any string you want to search for, whether in a field or not. It could be the name of a category or a flow. If you precede it with the optional `field:` the search
    - `operator` is one of the operators provided in the Apache Lucene search syntax.
- `<user>` is a user account that has the capability and permission to view and start a run of the flow.
- `<password>` is the password for the <user> account.

For more information on the Apache Lucene search syntax, see the Apache Software Foundation Web site's page on query syntax (http://lucene.apache.org/java/docs/queryparsersyntax.html).

Interrupting flow runs

Keep in mind the difference between interrupting and canceling 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 cancel a run, all information about the run is deleted. You cannot restart a canceled 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, 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 **Interrupt/Stop**.

   The run is stopped, and Central opens the Administration tab.

Reassigning ownership of a run

As an administrator (a member of Central's ADMINISTRATOR group), 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.

To reassign ownership of a flow

1. Logged in to Central with an account that is a member of the HP OO ADMINISTRATORS group, click the **Administration** 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.

	Name	Description	History ID | Run ID	Started ▲	Last Modified	User	State
☐	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 35 - 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 user account that you want to own the run, then click **Reassign**.

Resuming a run
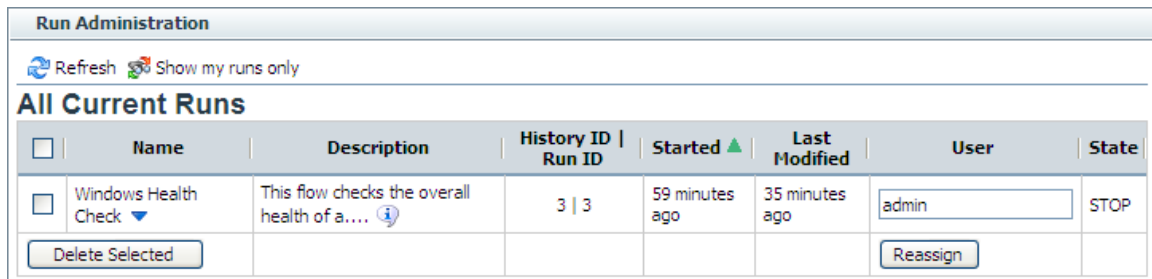
To resume a run that you did not start, you must be a member of the ADMINISTRATOR role. You resume runs on the **Administration** tab, which does not appear unless you are logged in with an account that has been added to that role.

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 is running 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.

Note: History ID and 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. On the other hand, various factors influence the assignment of Run ID, and Run ID is not universally unique across the history of the system. Within any given session, however, the Run ID should also be unique. When you click History ID|Run ID to sort the table of **Current Runs**, the table sorts on History ID.

To resume a flow run

1. Click the **Administration** tab.

Figure 27 - Run Administration

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

Figure 37 - Context menu for a flow run
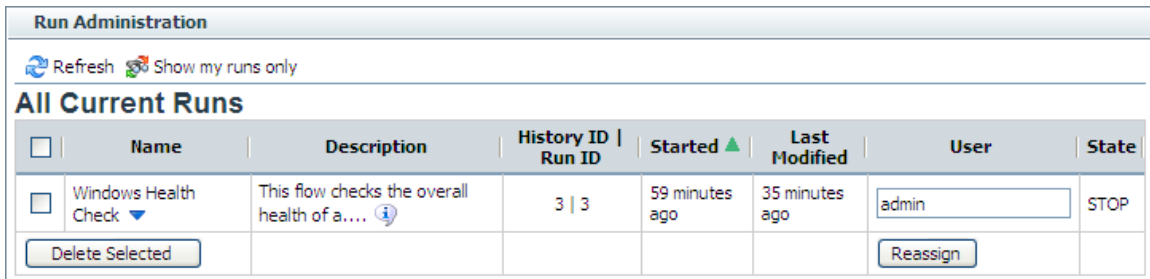
3. On the context menu that appears, click **Resume Run**.

Deleting a run

Note: If you delete a run from the Administration tab (which appears only if the logged-in user is a member of the ADMINISTRATOR group), History ID and 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. On the other hand, various factors influence the assignment of Run ID, and Run ID is not universally unique across the history of the system. Within any given session, however, the Run ID should also be unique. 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 **Administration** tab.

Figure 27 - Run Administration

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

Figure 39 - 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 27 - Run Administration

Note: If there is more than run current for a given 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 run

You can create a link to a flow run by capturing its URL, which you can place 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.

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. 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 **Link to this run**.

   The **Link to run** dialog box appears.
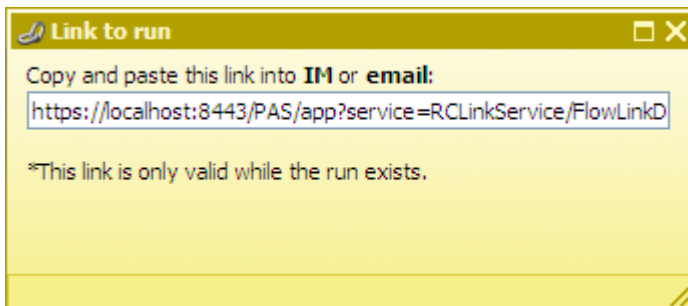
Figure 41 - Creating a link to the run

3. In the **Copy and paste this link...** text box, highlight the URL, and then close the **Link to run** dialog box.
4. In the external source from which you want to access the run, paste the shortcut that you copied.

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 access permissions that your account does not possess).

Note: The person who resumes the run must be logged in with an account that is a member of the HP OO ADMINISTRATOR role.

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

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 email message appears, with the URL of the flow included in the body of the message.
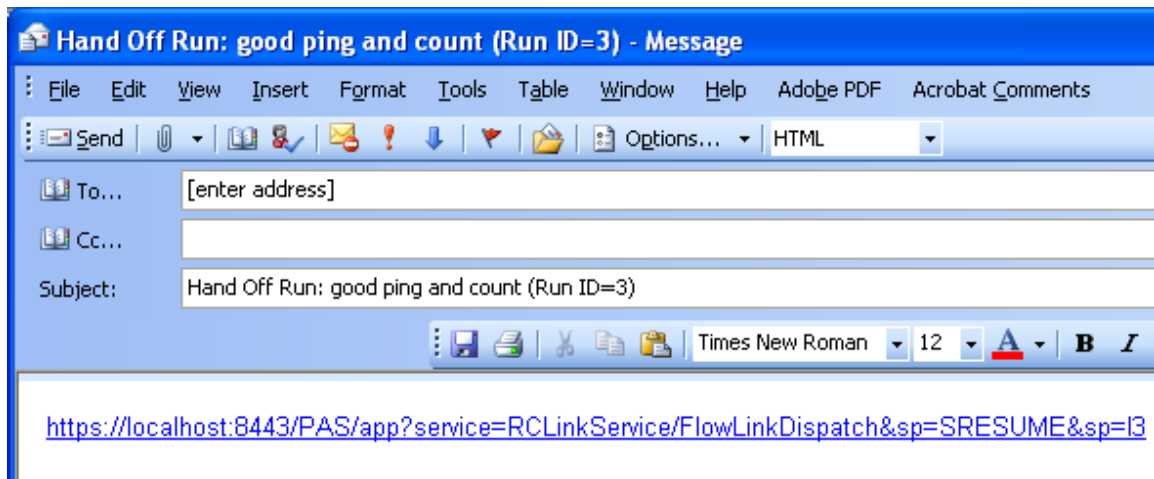
Figure 42 - Flow run handoff email message

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 email 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, HP OO administrators, and IT managers. For information on auditing and viewing reports on flow runs, see Run histories: What happened and why.

While users of Central can interrupt or cancel, resume or restart, flow runs that they own, administrative privileges are necessary to view all current flow runs and pause, resume, or delete them.

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

   Individual flow authors grant access permissions to their flows in studios. For information on assigning users or groups access permissions for various HP 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 HP OO for working with authentication by Active Directory, Lightweight Directory Access Protocol, or Kerberos providers and add HP 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 HP OO Administrator's Guide (AdminGuide.pdf).

Capabilities and access permissions

Working with flows, schedules, users, and other HP OO objects requires a combination of capabilities and (for flows and objects associated with them) access permissions that are particular to each object.

- A capability is the right to perform an action in HP OO, such as the MANAGE_USERS and MANAGE_GROUPS capabilities. AN HP OO administrator (a user with these just-mentioned capabilities) assigns groups the capabilities that they need. For more information, see Capabilities.
- Permissions are access rights to individual objects, such as individual folders, flows, operations, or system accounts. The four permissions are READ, WRITE, EXECUTE, and LINK, which flow authors grant to users or groups for individual objects. So:
  - To find and run a flow in Central, users must have read and execute permissions for the flow. In Studio, authors must have read, write, link, 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 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 HP OO.

Capability	Description
MANAGE_USERS	Create, delete, and modify internal users and map external user groups (groups that exist outside of HP OO) to internal HP OO groups. Only holders of this capability can create internal HP OO users.
MANAGE_GROUPS	Create, delete, and modify groups.
AUTHOR	Start Studio.
SCHEDULE	Schedule flows.
MANAGE_RUNS	View, delete, and reassign runs other than the user's.
RUN_REPORTS	Run reports and view metrics Dashboard pages.
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 HP 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 by 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.

HP 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		
	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 HP OO security, see the HP OO *Administration Guide* (AdminGuide.pdf).

When you first deployed HP OO, you mapped users to the HP OO groups. Depending on how you accomplish the mapping, when you deploy the HP 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 HP OO group. For information on mapping users to HP OO group, see the installation and deployment guide, *Installing Operations Orchestration software* (InstallGuide.pdf)*.*

Allowing external users into the Central system

Besides creating users within HP OO (called *internal* users), you can control who can use HP OO by mapping external users or groups to HP OO groups. Each internal user is created by hand on the **Administration** tab of Central, so large organizations will

authenticate user accounts for HP OO by enabling an external authentication provider—Active Directory (AD), Lightweight Directory Access Protocol (LDAP), or Kerberos—do the job. (For information on enabling authentication providers, see the next section.)

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

- The external user's account is a member of an external group that is mapped to an HP 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.

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.

To enable HP OO for one or more authentication providers

1. To enable authentication provider, select the appropriate check box (**AD Enabled**, **LDAP Enabled**, or **Kerberos Enabled**).

2. Modify the configuration values for the authentication provider according to your organization's needs, according to one of the sections that follow this procedure:
   - AD authentication settings
   - LDAP authentication settings
   - Kerberos authentication settings

3. After configuring settings according to one of the sections listed in the preceding step, test the current settings for an authentication provider from within HP OO by clicking the **Test** button for the authentication provider.

   The **Testing AD Settings** dialog box appears.
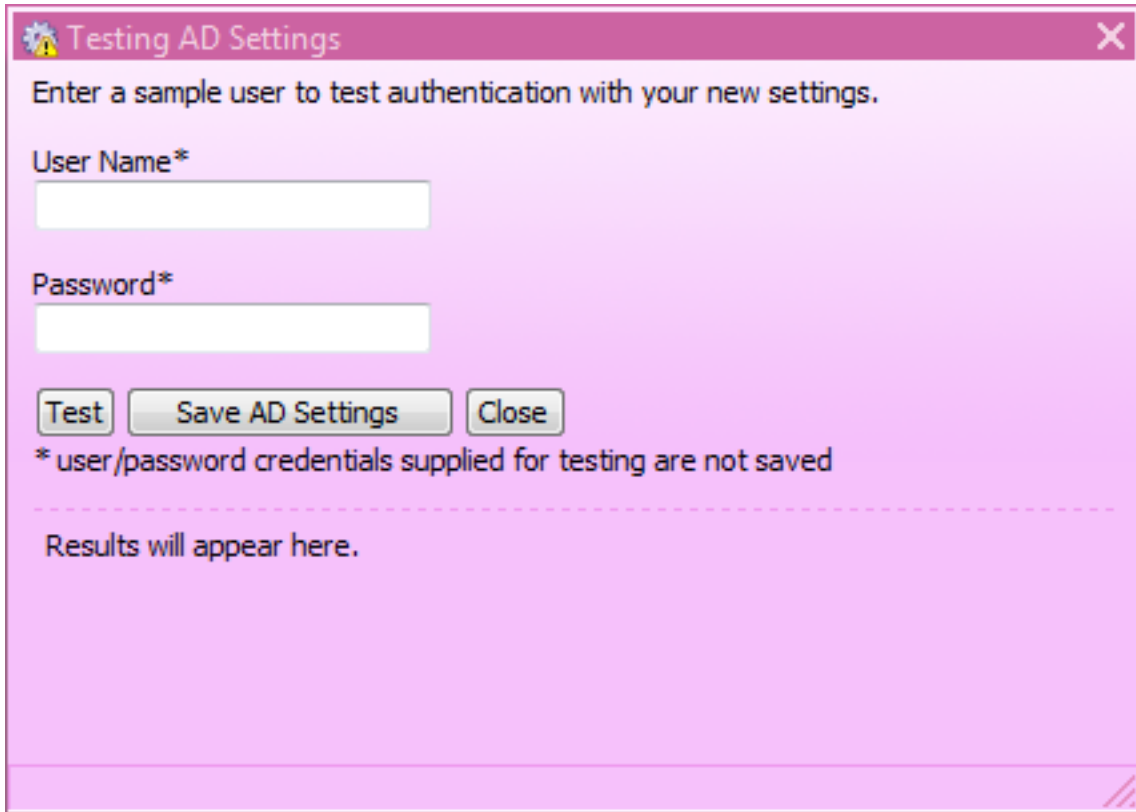
Figure 29 - Testing an authentication provider

4. Type the user name and password of an external account that is authenticated by the provider you're testing, and then click **Test**.

5. If the test fails, modify the settings and test again.

6. When the authentication provider settings test successfully, click the **Save AD Settings** button for the authentication provider.

7. When you see a message to restart the RSCentral service, do so.

8. After configuring the system,

For the following sections on configuring HP OO using AD or LDAP 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.

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
	Save AD Settings Test Refresh

Figure 30 – AD Authentication Settings

To authenticate with Active Directory

1. Select the **AD Enabled** check box.

2. In the **Value** box for the **default role**, specify which HP 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**.

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 HP OO:

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

   Next we need to create search filters, to tell HP 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 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 HP 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 HP OO *Administration 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.

Note that the backslash in the domain\user syntax here is rendered as a double backslash:

`<domain>\\{0}`

Be sure to type **{0}** exactly as it appears above.

8. To specify the contexts in which HP 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 ☑ 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
	Save AD Settings Test Refresh

Figure 31 - AD authentication enabled

9. To save your settings, click **Save AD Settings**.

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 HP OO groups, see Mapping external groups to HP 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.

Figure 32 - Settings for configuring HP 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 HP 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**.

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 HP OO:

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

Next we need to create search filters, to tell HP 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 HP 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 HP OO *Administration 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 HP 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 HP OO group names: In the **Value** box for the **List of LDAP contexts containing 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**.

If your IT organization also authenticates with Active Directory and/or Kerberos, complete the procedures in <u>AD authentication settings</u> or <u>Kerberos authentication settings</u>. Finally, to map external LDAP groups to HP OO groups, see <u>Mapping external groups to HP OO groups</u>.

Kerberos authentication settings

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

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

Kerberos Authentication Settings ☐ 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.	
KDC host (format is "address:port" or "address" when using the default port).	MyCompany.ad
Kerberos realm.	MYCOMPANY.AD
The default group an Kerberos authenticated user gets when there is no explicit group assigned.	EVERYBODY
	[Save Kerberos Settings] [Test] [Refresh]

Figure 33 - Settings for configuring HP 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 HP OO home directory, and the path should be relative to that directory.

   For instance, the example in the description to the left of the **Value** box describes the Kerberos configuration file as being in the \Central\conf subdirectory of the HP OO home directory.

   Important! If the specified file path for the Kerberos configuration file is invalid, HP 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 HP 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**.

6. To save your settings, click **Save Kerberos Settings**.

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 HP OO groups. For more information, see [anaging users](#).

Managing users

Users are either **internal** to HP OO—that is, you create them within HP OO and they do not exist outside of HP OO—or are **external** user accounts that exist independently of HP 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 HP 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 HP OO groups, or map the user's external (AD or LDAP) account to an HP OO account.

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 HP 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 HP OO users for those authors to log into Central with when testing their flows.
- The production server, however, might have two dozen or so IT personnel logging into Central in order to run flows, in addition to administrators and managers who might need to log 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 HP OO groups.

Note: If an external group has the same name as an internal HP OO group (after translation of the external name using HP OO group-name rules), then members of the external group can log in to HP OO with the capabilities of the HP OO group. HP 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 HP OO with the capabilities of the HP 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 HP OO.

Adding a user

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.
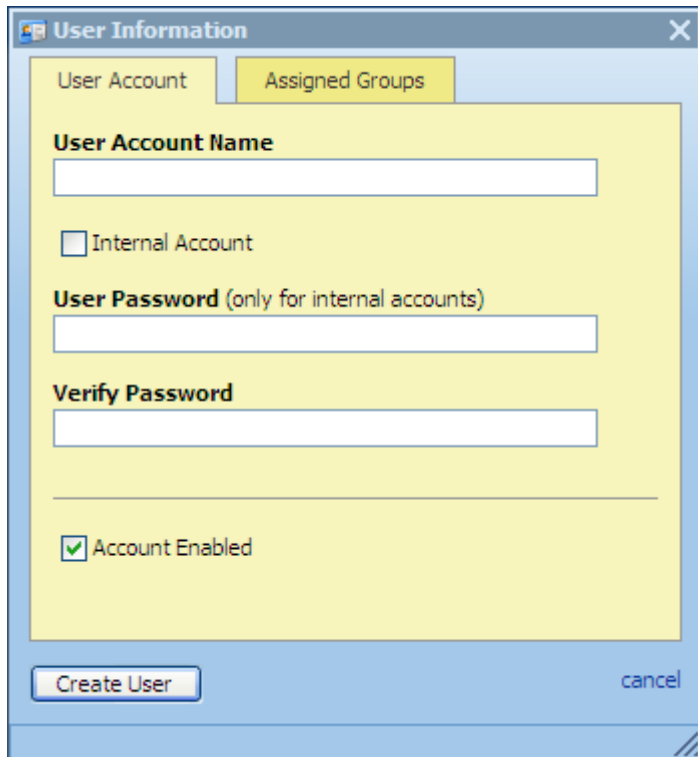
Figure 34 - User Information dialog box

2. Type the user account name.

3. If you are creating the account within HP OO, select the **HP OO Internal Account** check box, and then type and verify a password for the user to log in to HP OO with.

   The password must be at least six characters long.

   Note that by default, the **Account Enabled** checkbox 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 *Hewlett-Packard Software Operations Orchestration Administration Guide* (AdminGuide.pdf)*.*
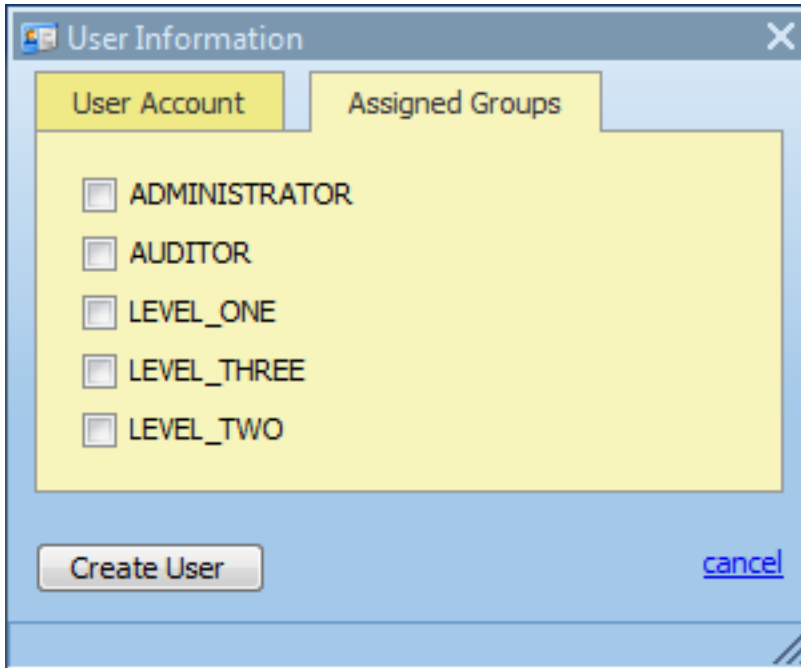
Figure 35 - Assigning groups to a user

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 HP OO even if all other accounts are disabled or otherwise nonfunctional.

To make changes to a user's account

1. On the **Administration** tab's **Manage Users** subtab, click the notepad-and-pencil icon ( ), then, in the **User Information** dialog, make changes in the same way that you configured the account when you added the user.

2. When you've finished making changes, click **Update User**.

Deleting a user

To delete a user account

- On the Manage Users subtab of the Administration tab, select the checkbox 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.

To make changes, you must be a member of a group that has the MANAGE_GROUPS capability.

HP OO groups make it easy to add users as groups to HP OO and to manage their capabilities and rights. You can map your IT organization's AD or LDAP groups to HP OO groups, thus adding the entire membership of the group at once, as HP OO users.

Scenario

Suppose you want to map the following groups in your IT organization to existing HP 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 HP 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 HP OO groups and create rights for the HP 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.

  The following are three special HP OO groups:

- 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 HP 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 HP OO automatically becomes a member of this group. The group doesn't have any capabilities, but does have access to certain HP OO objects, such as Accelerator Packs. As a result, the HP 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.
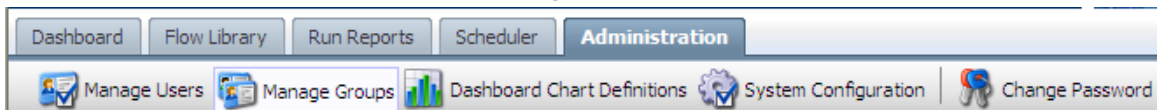

Figure 36 - Manage Groups subtab

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

- Mapping external groups to a group.
- Changing a group's capabilities.
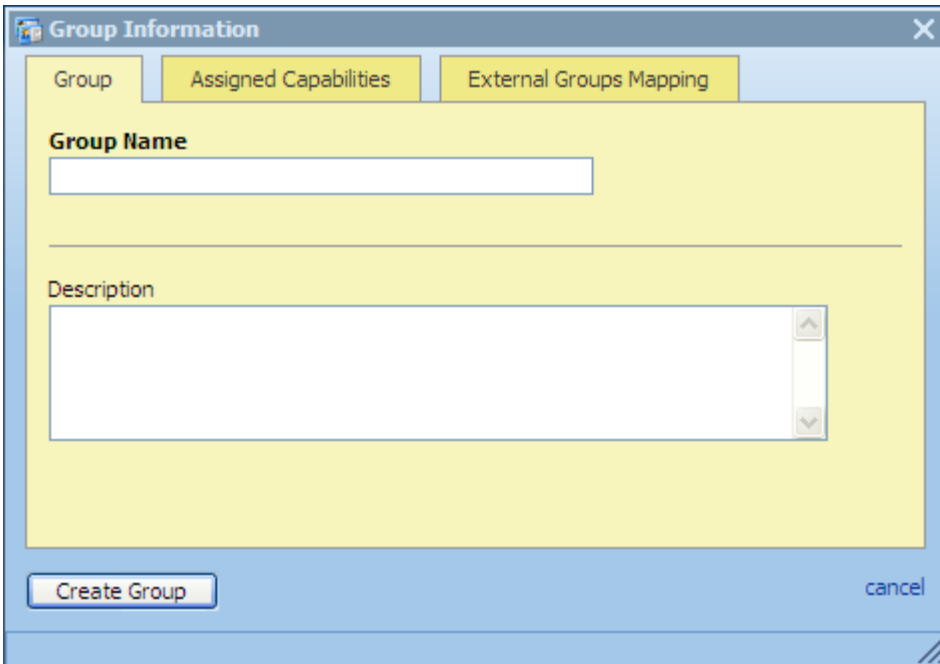- Changing the group's description or name.

Adding groups

To add a group

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

Figure 51 - Administration tab

2. Under **Groups**, click Add **New Group**.

   The **Group Information** dialog box appears.

Figure 52 - 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 **Update 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 HP OO groups. For more information on using external groups, see the HP OO Administrator's Guide (AdminGuide.pdf).

Adding HP OO users to groups

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

To assign a user to an HP 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 **Groups** tab, and then specify the groups you want to add the user to.

# Mapping external groups to HP OO groups

To add users of the following AD or LDAP groups to HP OO groups as shown in this table, use the **Manage Groups** subtab on the **Administration** tab.
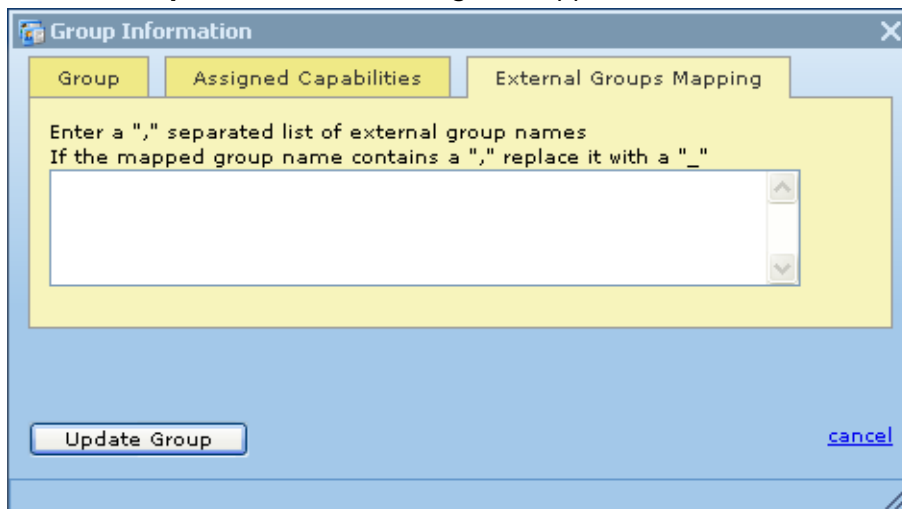
AD or LDAP group	HP OO group
Service Desk	LEVEL_ONE
Network Specialist	LEVEL_THREE
Manager	ADMINISTRATOR

Note: If an external group has the same name as an internal HP OO group (after translation of the external name using HP OO group-name rules), then members of the external group can log in to HP OO with the capabilities of the HP OO group. HP 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 HP OO with the capabilities of the HP 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 HP OO.

To map external groups to HP OO groups

1. On the **Administration** tab, click the **Manage Groups** subtab, and then click the edit icon (   ) in the row of the group you want to map the external group to.

2. In the **Group Information** dialog that appears, click the **External Groups** tab.

Figure 37 — 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 HP OO group.

   For instance, to map the AD or LDAP "Network Specialist" group to the HP OO LEVEL_THREE group, type **Network Specialist** in the text box, and then click **Update Group**.

   After making the above group assignments, the

You create and manage HP OO user accounts, manage group membership, and assign capabilities (defined actions) on the **Administration** tab of Central. For more on HP OO groups and capabilities, see the Administration 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 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

The ability to do things with flows and the objects associated with them

To change the capabilities assigned to a group

- In the **Group Information** dialog, click the **Assigned Capabilities** tab.
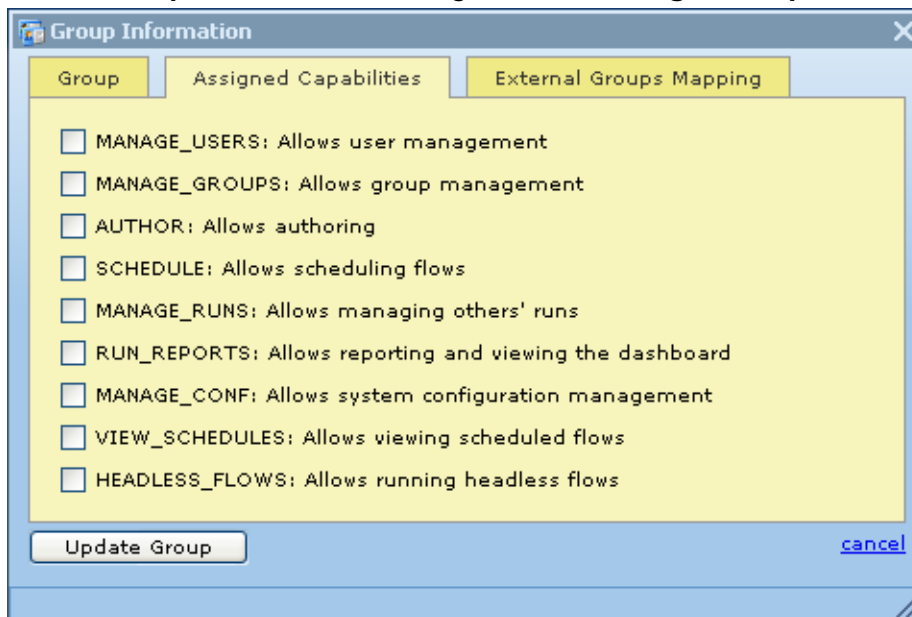
Figure 38 - Assigning capabilities for a group

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 *Operations Orchestration Administration 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 *Operations Orchestration Administration Guide.*

- 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.
- 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 Administration Guide.*

- Configuring Active Directory to run over SSL.
- Configuring HP 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 HP 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.

To delete a group

1. On the **Administration** tab, click **Manage Groups**.
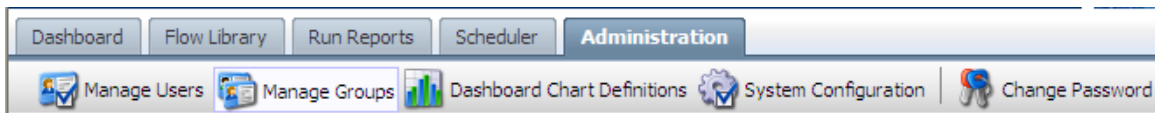
Figure 55 - Administration tab

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

Figure 56 - 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 **Administration** page, 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

To see which flows are currently running

1. Log on to Central with an account that has HP OO administrative permissions.
2. Click the **Administration** navigation tab.

   OR

   If you have been working on one of the subtabs of the **Administration** tab, click the **Administration** tab again.

   The **Run Administration** area appears, and the table displays all current flow runs, including the following information.

	Name	Description	History ID | Run ID	Started ▲	Last Modified	User	State
☐	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 39 - 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.

3. To refresh the page, click **Refresh.**
4. To show only your own runs, under Run Administration and above All Current Runs, click **Show my runs only**.

   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:

   - On resuming a run, see Resuming a run.
   - On deleting a run, see Deleting a run.
   - On viewing a run's history, see Run histories: What happened and why.

Other system configurations

There are several other HP OO system configurations that you can change:

- Return on Investment (ROI) reporting
- How frequently the Dashboard charts refresh
- LDAP referrals
- Ability to directly connect to a public repository
- Configuring the HP OO server process heap size.

To make these changes, you must be a member of a group that has the MANAGE_CONF capability, which enables you to make changes to the system configuration.

Enabling ROI reporting

By default, ROI is reported in:

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

However, you can disable ROI reporting.
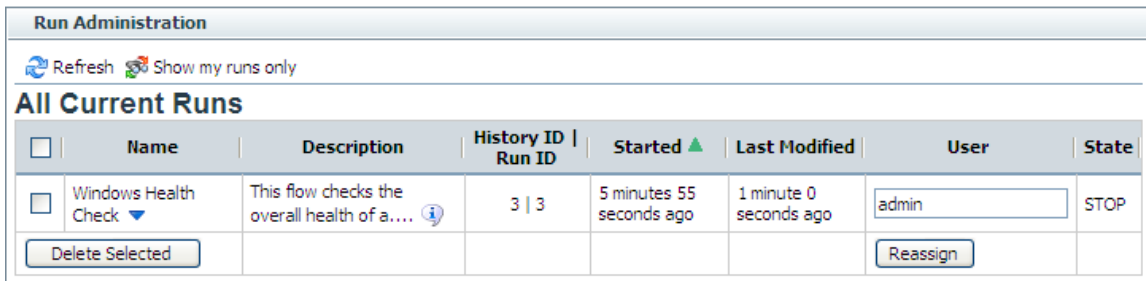
To enable or disable ROI reporting

1. Log on to Central with an account that has HP OO administrative permissions.
2. Click the **Administration** tab, and then the **System Configuration** tab.
3. To enable ROI reporting, in the **General Settings** area:
   - In the **Enables ROI** row, type **true** in the **Value** box.
   OR
   - To disable ROI reporting, type **false** in the **Value** box.
4. Click **Save General Settings**.
5. Restart the HP 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.

To change the rate at which Dashboard charts refresh their data

1. Log on to Central with an account that has HP 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.
4. Click **Save General Settings**.
5. Restart the HP OO Central service (RSCentral).

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.

To specify how Central authentication manages LDAP referrals

1. Log on to Central with an account that has HP 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.
4. Click **Save General Settings**.
5. Restart the HP 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.

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

1. Log on to Central with an account that has HP 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 HP OO Central service (RSCentral).

Controlling which authors can directly connect to the shared Central repository

By default, only authors who have administrative privileges (that is, are members of the ADMINISTRATOR group) can directly connect to the shared Central repository.

Directly connecting to the shared Central repository means opening the shared Central repository in Studio and making changes to it.

It is strongly recommended that you leave this default setting as is. Allowing non-ADMINISTRATOR authors to make changes to the public repository has two effects:

- Creates a security vulnerability for HP OO and flows.
- Reduces control over the changing of flows used in your production environment.

To enable non-ADMINISTRATOR authors to directly connect to the public repository

1. Log on to Central with an account that has HP OO administrative permissions.
2. On the **Central Administration** tab, click the **System Configuration** subtab.
3. In the **Value** box for the **Only administrative users can connect directly...** setting, change **true** to **false**.
4. Click **Save General Settings**.
5. Restart the HP OO Central service (RSCentral).

Changing the configuration of a Central cluster

For consistency in the clustering nodes, certain changes that you make in one node of the cluster must be reflected in the cluster's other nodes.

For more information on Central failover and run-recovery clusters, see the *HP OO Installation Guide*.

To change configuration of Central nodes in a cluster for failover and run recovery

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

Clustering Settings ☐ Clustering Enabled

Description	Value
The class D multicast address to use when UDP is selected as the protocol.	228.10.10.10
The multicast port to use when UDP is selected as the protocol.	45566
The name of the cluster this Central belongs to.	CENTRAL_CLUSTER
The protocol to use to communicate with other Central cluster members (UDP or TCP). UDP specifies to use UDP IP multicasting between cluster nodes. TCP specifies to use TCP communication between cluster nodes.	UDP
A comma separated list of host[port] specifications of all Central hosts and their TCP listener ports in the cluster (i.e. hosta.mycompany.com [45566],hostb.mycompany.com[45567]). This is only used if TCP is selected as the protocol.	localhost[45566]
The TCP listener port to use when TCP is selected as the protocol.	45566
	Save Clustering Settings Refresh

Figure 58 - Creating and configuring a failover and run recovery cluster

Each Central cluster requires a unique pair of settings for the cluster's multicast address (mcast_addr) and the multicast port (mcast_port). The values used for the mcast_addr and mcast_port must not overlap with the IP address/port combination used for other applications that run in the environment.

3. In the **Value** box for **The class D multicast address to use when UDP is selected as the protocol** setting, type the IP address to use for multicasting.

   The address you specify must not to conflict with other applications

4. In the **Value** box for **The multicast port to use when UDP is selected as the protocol** setting, either retain the default (45566) or change it to a port number that:

   • Is not blocked on any of the cluster nodes.

   • Is not used for other purposes in your production environment.

5. In the **Value** box for **The name of the cluster** setting, name the cluster.

   The cluster name must be consistent across the cluster nodes.

6. In the **Value** box for **The protocol to use to communicate with other Central cluster members...** setting, select either UDP or TCP.

   The protocol used must be consistent across the cluster nodes.

7. If TCP is specified for the cluster's internal communication protocol, in the **Value** box for the **A comma-separated list of host[port] specifications of all Central hosts...** setting, type the name-port pairs of the cluster nodes, enclosing each port number in square brackets and separating the pairs with commas.

   The list of cluster nodes must be consistent on each of the cluster's nodes.

   For instance, if your nodes were:

   • edgar.mydomain.ad, using port 888

   • rosalind.mydomain.ad, using port 555

   Important: The port numbers that you specify must not be blocked on any of the cluster nodes, nor be used for other purposes in your production environment.

   You would type the following in the **Value** box:

   edgar.mydomain.ad[888],rosalind.mydomain.ad[555]

8. If TCP is specified for the cluster's internal communication protocol, in the **Value** box for the **The TCP listener port to use...** setting, type the port over which the Central node will listen to the other cluster nodes.

9. Click **Save Clustering Settings**.

For information on changing the network address binding on Central servers that have more than one network interface, see the HP OO *Installation Guide* or *Administration Guide.*

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 HP 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 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.

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.