

HP Operations Orchestration

For the Windows ®, Solaris ®, and Linux operating systems

Software Version: 9.05

Central User Guide

Document Release Date: June 2012

Software Release Date: June 2012



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 Notice

© Copyright 2005-2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

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

UNIX® is a registered trademark of The Open Group.

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

Documentation Updates

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

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

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

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

Support

Visit the HP Software Support Online web site at:

<http://www.hp.com/go/hpsoftwaresupport>

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

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

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

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Disclaimer for PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format.

Note: Some topics do not convert properly to PDF, causing format problems. Some elements of online help are completely removed from the PDF version. Those problem topics can be successfully printed from within the online help.

Contents

Central User Guide.....	1
Contents.....	6
Quick View of Operations Orchestration Central.....	10
A scenario.....	10
Starting Central.....	12
Central Web application: the initial page.....	13
Flow Metrics area.....	14
Popular flows.....	15
Navigating in Central.....	16
Changing your OO Central password.....	17
Viewing the OO groups you belong to.....	18
Finding flows.....	19
Browsing flows in the Library.....	19
Searching for a flow.....	20
Types of flow runs.....	23
Running a flow.....	25
Restarting a flow.....	27
Running subflows.....	27
Previewing flows.....	29
Navigating in the flow preview.....	31
Deleting a run from within the run.....	33
Updating a run.....	34
Opening the flow graph in a separate browser window.....	35
Seeing what has happened in a flow run.....	37
Run histories: What happened and why.....	41
Exporting run history reports.....	44

Run histories in more detail	45
Viewing step results and other advanced data	47
Scheduling flow runs	49
Run-scheduling concurrency	49
Creating a schedule for a flow	50
Editing flow schedules	54
Enabling and disabling flow schedules	55
Deleting flow schedules	56
Configuring Scheduler settings	57
The Dashboard: Learning more from flows	59
Collating data on Dashboard reporting charts	59
What do the bars tell you?	61
Learning more from the charts	62
Creating and modifying charts	64
Exporting and importing chart definitions	67
Starting a flow from outside Central	69
Creating a link to a flow run	69
Creating a link to a flow	70
Handing off flow runs	71
Auditing and managing flows	72
Users, groups, and access control	73
Capabilities and access permissions	73
Capabilities	74
Permissions	74
Allowing external users into the Central system	76
Using external authentication for Central users	77
Example data	77
AD authentication settings (Deprecated)	77
LDAP authentication settings	81
Kerberos authentication settings	86
Managing users	89
Internal or external users?	89

Adding a user.....	90
Editing a user's account.....	91
Deleting a user.....	92
Managing groups.....	93
Scenario.....	93
Default groups on the Manage Groups subtab.....	95
Adding groups.....	95
Adding OO users to groups.....	96
Mapping external groups to OO groups.....	97
Changing a group's assigned capabilities and description.....	98
Deleting groups.....	100
Managing flow runs.....	101
Refreshing the Current Runs tab.....	102
Filtering the flow runs to view.....	102
Filtering flow runs using the Filter Panel.....	102
Filtering flow runs with column headers.....	105
Resetting filters to their defaults.....	106
Saving filters.....	106
Loading a saved filter.....	107
Deleting a saved filter.....	107
Viewing the details of a flow run.....	108
Taking control of a flow run.....	110
Resuming a flow run.....	111
Deleting flow runs.....	111
Reassigning ownership of flow runs.....	112
Watching runs.....	114
Viewing flow run properties.....	115
Managing Central nodes.....	116
Viewing the Central nodes.....	117
Taking nodes offline.....	117
Bringing nodes online.....	118
Integrating a load balancer with node management.....	119

Viewing RAS connections.....	121
Managing your runtime environment with overrides.....	122
Adding a RAS override.....	124
Adding a system account override.....	126
Adding a system property override.....	127
Editing environment overrides.....	128
Deleting environment overrides.....	130
Best practices for runtime environment overrides.....	130
Other system configurations.....	131
Enabling the visibility of ROI reporting.....	132
Restricting authoring to administrative users.....	133
Setting how many versions of an object are stored.....	133
Restricting who can publish.....	134
Changing the Dashboard charts refresh rate.....	135
Specifying a required prefix for init params when using a URL to start a flow run in.....	
Central.....	135
Specifying how Central manages LDAP referrals.....	136
Enabling Central to resume runs that were interrupted when Central failed.....	136
Enabling load balancer management for a Central cluster.....	137
Troubleshooting.....	139
Web browser shows a security-certificate-related warning when you open the Central....	
Web site.....	139
I was sent back to the login page.....	139
I cannot restart or resume a flow that I started or that was handed off to me.....	139
I cannot change or create a schedule for a flow.....	139
Central fails to create a schedule for a flow.....	140
Changes were made to the Public repository, but they don't appear in the Flow Library....	140
Flows run under heavy load with command-line or RAS operations fail on Windows.....	
systems with error code 128.....	140

Quick View of Operations Orchestration Central

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

You access and run flows from Central to:

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

This Quick View of Central will show you how Central:

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

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

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

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

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

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

A scenario

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

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

- Central users run the flows that resolve the alerts.

In the Flow Library (on the **Flow Library** tab), the OO Central user locates and runs the flows that resolve the situations causing the alerts.

- In the Dashboard (on the **Dashboard** tab), users then analyze the data that is produced by the flows that resolve the alerts to discover information such as:

- What are the alerts that are showing up most frequently?
- What is the outcome for each alert?
- Which server or application generated the most alerts?
- Which flows ran most often, and what were their outcomes?
- Which applications and servers had fatal errors?
- How many alerts of various severities were there?

Starting Central

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

Central provides a graphical user interface for:

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

Before starting Central, make sure that the server on which the shared repository resides is up and running.

To start Central

1. Start your machine's Web browser.
2. Paste the URL that your administrator sent you into the **Address** box of the Web browser and then press Enter.
3. When the message appears that you are about to view pages over a secure connection, click **OK**.

If a message appears, warning you that the site is not trusted, it is safe to proceed.

4. Click **Yes**.
5. When the Central Login page appears:
 - a. Type your user name and password.
 - b. Click **Login**.


Central opens, and you're ready to locate, run, and view information on flows.

Note: If you receive an error message when logging on to Central, check to make sure that

JavaScript is enabled on your machine.

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.



Notes:

- Whether you see the **Current Runs**, **Reports**, or **Scheduler** navigation tabs depends on which capabilities (the abilities to schedule flows, to manage other users' runs, etc.) are assigned to your group or groups.
- The time given in the title bar of OO Central is in the timezone where the client machine on which you're running OO Central resides.

The title bar displays:

- The relationship between the Central client's time zone and Greenwich Mean Time (GMT).
All the times shown in the Central web application, such as in the **Scheduler** and **Reports** tabs, are the Central client's time, and thus have the same relationship to GMT as that described in the title bar.
- Your user name.
- The first group to which you belong (from an alphabetical list of groups to which you belong).
For information on viewing the groups you belong to, see [Viewing the OO groups you belong to](#).

- A link to the OO Central Help system.
- The **Logout** button.


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

Flow Metrics area

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

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

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



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 (beside the graph title).

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.

Flow Utilization for the Last 7 Days



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

Popular flows

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

▼ Popular Flows*				
Most-executed Flows for the Last 7 Days				
Flow Name	Number of Times Run	Average Run Time [secs]	Flow ROI Value	Start Run All
Windows Health Check	1	43.297	.00	
Simple SMTP Check	1	16.172	.00	
Restart Service - Tutorial Flow	1	19.749	2.15	

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


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

To open a Dashboard chart

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

Navigating in Central

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



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

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


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

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

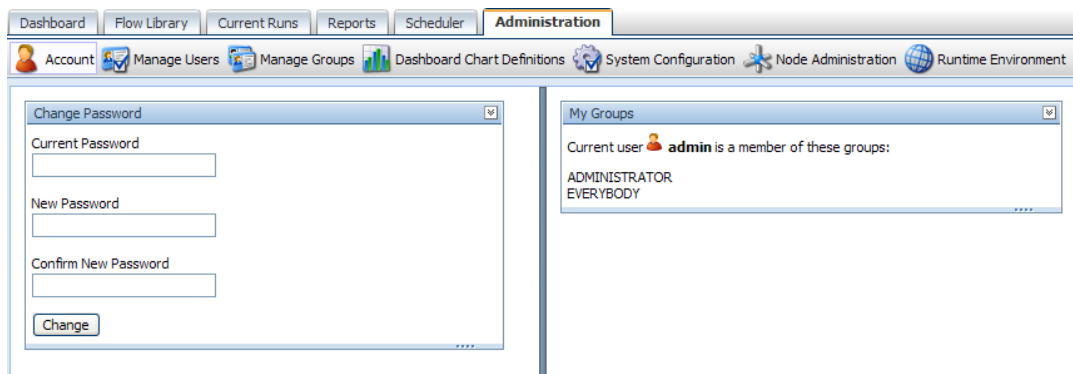
Changing your OO Central password


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

To change the password for the current user

1. Click the **Administration** tab, and then click **Account** ( **Account**).

The **Account** tab opens.



The screenshot shows the OO Central Administration interface. The top navigation bar includes tabs for Dashboard, Flow Library, Current Runs, Reports, Scheduler, and Administration (which is selected). Below the navigation bar, there are several sub-tabs: Account (selected), Manage Users, Manage Groups, Dashboard Chart Definitions, System Configuration, Node Administration, and Runtime Environment. The main content area is divided into two panels. The left panel, titled 'Change Password', contains three text input fields labeled 'Current Password', 'New Password', and 'Confirm New Password', along with a 'Change' button. The right panel, titled 'My Groups', displays the text 'Current user  admin is a member of these groups:' followed by a list of groups: ADMINISTRATOR and EVERYBODY.

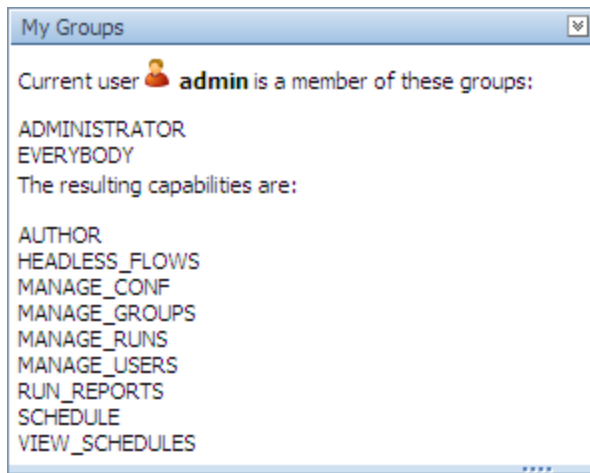
2. In the **Current Password** text box, type your current password.
3. In the **New Password** and **Confirm New Password** text boxes, type the new password, and then click **Change**.

Viewing the OO groups you belong to

To view the group memberships of the current user:


- Click the **Administration** tab, and then click the **Account** button  Account

The **My Groups** panel lists the groups to which the logged-in user account belongs.



Finding flows

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




The OO Library

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

If you can't find the flow you need, you can [search](#) for it.

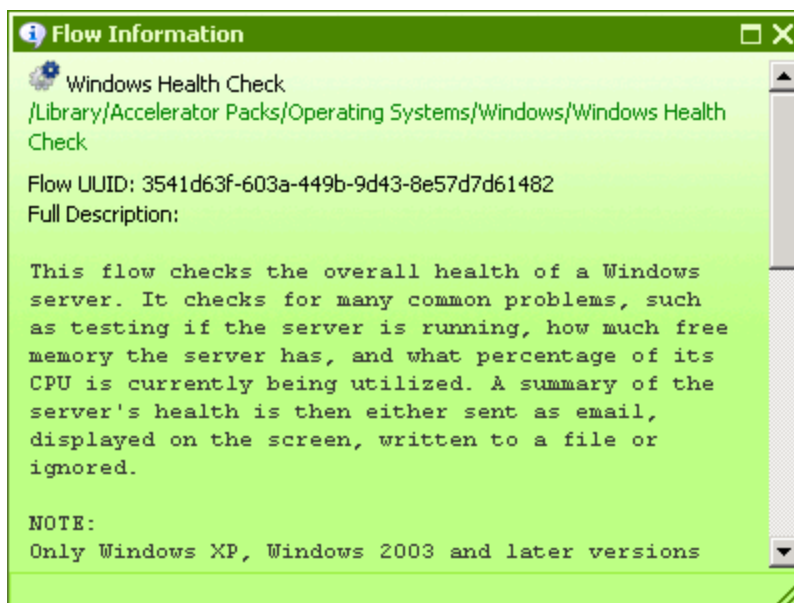
Browsing flows in the Library

To browse the Library for a flow

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

3. To see detailed information about a flow, right-click the flow name, and then click **About**.

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



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

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.

There are three ways to run a flow:

- **Guided Run**
You click to carry out each step and respond to any user prompts.
- **Run All**
The flow completes all the steps on its own, and you only respond to user prompts.
- **Instant Run**
Like Run All, except that the user prompts, results, and other data generated appear in a dialog box.

For more information about running a flow, see [Types of flow runs](#) and [Running a flow](#).

Searching for a flow

Central's search mechanism uses the Apache Lucene search syntax. In addition to the basic search method described in this topic's procedure, you can use the syntax to construct more highly targeted searches. For more information on the search syntax, see the Apache Software Foundation Web site's page on query syntax.

To search for a flow

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

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

OR

Type a search field and value using the form:

```
<fieldname>:<term>
```

where `<term>` is the particular value in the field that may find the desired flow. For example, to search for flows that have a step named Notify in them, enter:

```
stepname:Notify
```

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.



The screenshot shows a search results page with a search bar containing 'stepname:Notify' and a 'Search' button. The results are displayed in a list format, with each result including a gear icon, a title 'Check Connectivity', a path, a description, and a list of inputs and responses. The first result is located at '/Library/ITIL/Incident Management/Network/Check Connectivity' and the second is at '/Library/ITIL/Incident Management/Servers/Check Connectivity'. The search results are paginated, with '<Previous Page | Next Page>' links.

```
Search Results
You searched for the text:stepname:Notify

stepname:Notify Search

<Previous Page | Next Page>

Check Connectivity
/Library/ITIL/Incident Management/Network/Check Connectivity
Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable.

Inputs:
Host - The host to ping from.
username - the username for host.
password - the password for host.
target - the host to ping.
notifyMethod - How you want the report delivered.

Results:
report - Connectivity test results

Responses:
success - The test passed.
diagnosed - A problem was found
failure - something went wrong, or no response was received for the ping.

-----

Check Connectivity
/Library/ITIL/Incident Management/Servers/Check Connectivity
Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable.

Inputs:
Host - The host to ping from.
username - the username for host.
```

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.

Types of flow runs

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 

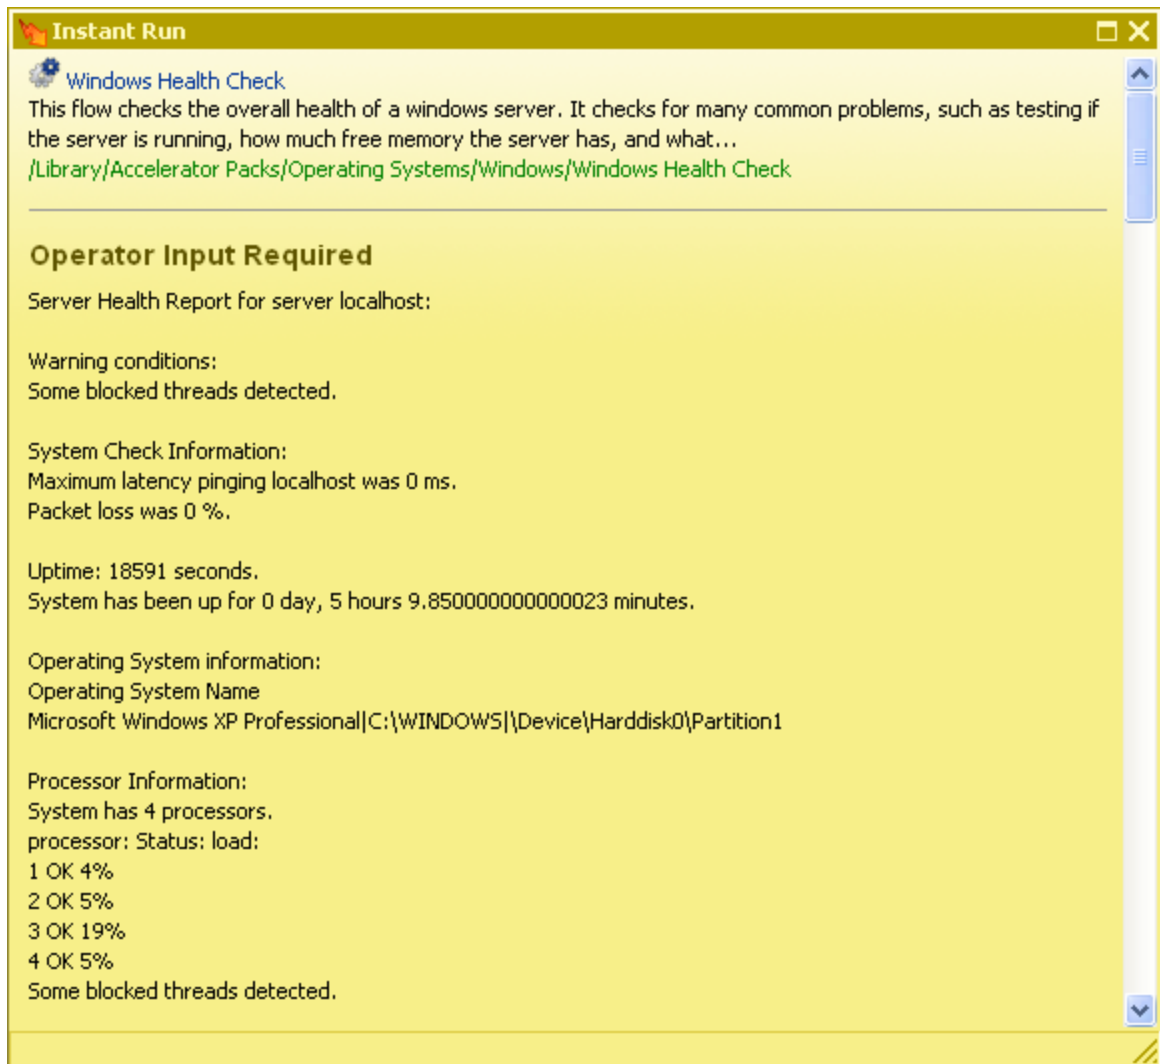
In run all, when you start the flow, it runs straight through to completion, except for any user-prompt inputs that require you to supply a value.

- Instant run 

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

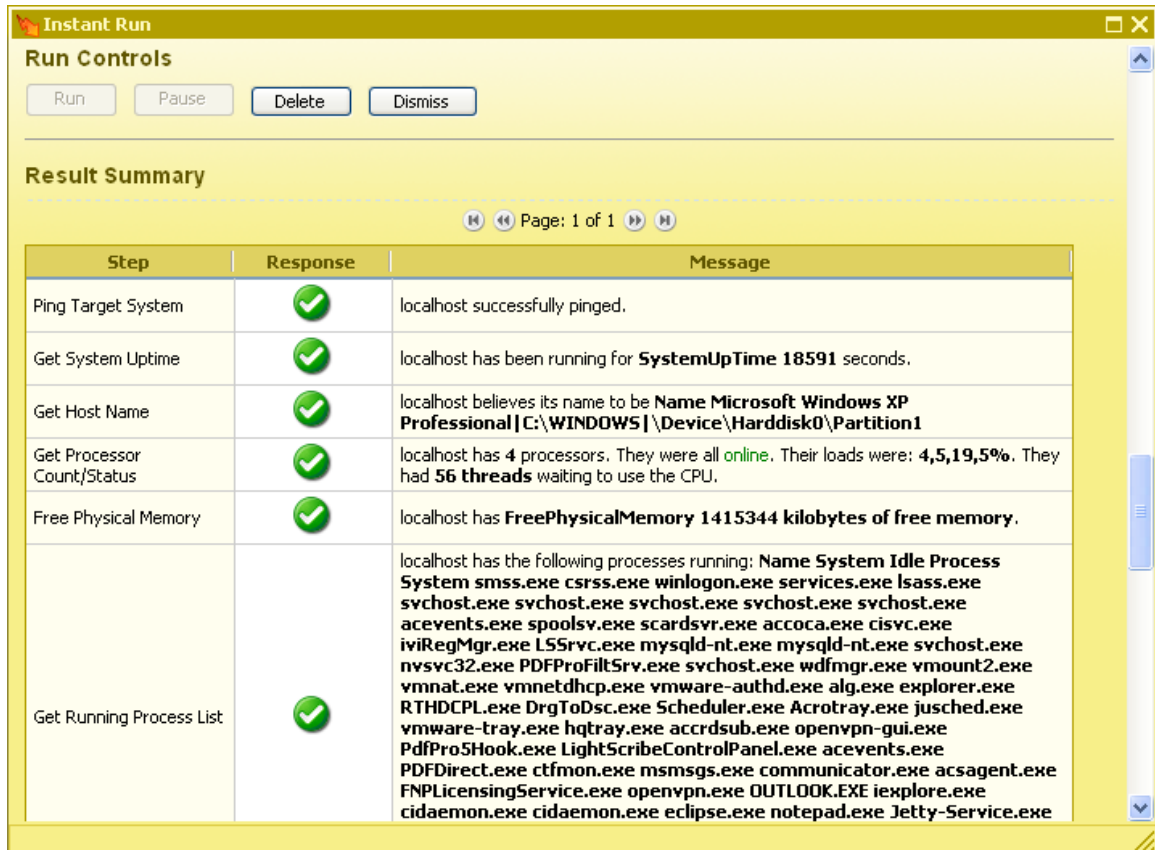


Only the essential controls for completing the run are available, and only while the box is open. Note though, the information that is visible when you expand the **Instant Run** dialog box.



Note: If your account does not have sufficient permissions or the capability to start a run of this flow, a message appears, alerting you to the fact and providing a button for handing off the flow. For required inputs (such as credentials) for which you need to provide values, text boxes appear in the **Instant Run** dialog box.

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



For information on the ways you can start a flow run from within Central, see [Ways to run flows](#). 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](#).

Running a flow

There are several ways to start a flow run. The first two are described in this topic.

- From the **Flow Library** tab
- From the **Flow preview** page
- By scheduling one or more runs of a flow

For information on scheduling flow runs, see [Scheduling flow runs](#) and [Creating a schedule for a flow](#).

- From outside the OO Central web application, using a command-line command

For information on starting a flow run from outside Central, see [Starting a flow from outside Central](#).

To run a flow directly from the Flow Library page

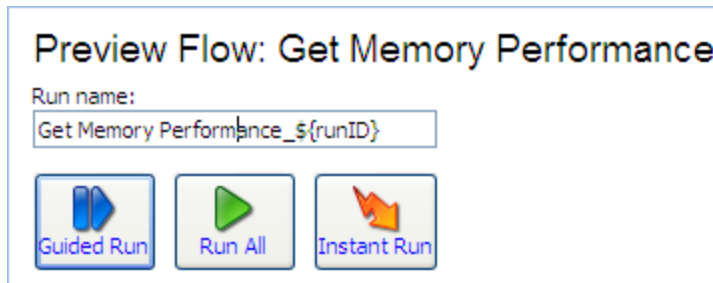
1. On the **Flow Library** tab, navigate to the flow you want to run.
2. Right-click the flow name, and then click one of the following:

- **Guided Run**
- **Run All**
- **Instant Run**

To run a flow from the Preview Flow page

1. On the **Flow Library** tab, click the flow name to open the flow's **Flow Preview** page.

The **Flow Preview** page includes, in the top-left pane, a button for each type of run and the **Run name** box, containing the default name of the run.



2. To change the run name from the default, type a new name in the **Run name** box.

The default run name is made up of the name of the run and the number of the run, which is obtained dynamically by the variable reference `${runID}`. The run name will be something like "Windows Health Check_2". You can create a different run name, if you wish.

Notes on the run name:

- The run name is limited to 255 characters.
- When you change the run name, the run's URLs in the **Execution Links** panel on the flow's preview page are automatically also changed.
- When you manually run a flow by right-clicking the flow name in the Flow Library and then selecting one of the run types, the run name is given the default run name.
- You can specify a run name on the **Flow Preview** page, in the command with which you start a run from outside Central, or when you create a schedule of runs of the flow. By default, all the runs in the schedule are given the schedule's name.

3. Click one of the run-type buttons:

- **Guided Run**
- **Run All**
- **Instant Run**

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

- **Results Summary** panel: shows the step names, their responses, and what happened during each step
- **History Tree** panel: shows the order in which the flow steps and subflow steps were run and the results of the steps

- **Context** panel: shows all of the flow variables used in the flow run and the values assigned to them
- **Stack** panel: shows the root flow and, if it uses a nested-execution pattern, which subflow you may currently be inside
- **Graph** panel: graphic depiction of the flow that highlights each step as it is executing

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

Restarting a flow

Anywhere except in the Central display of a flow that is running, restarting a flow is the same as starting a flow run. After a flow's guided run or run-all run completes, the Central display of the flow's run state includes the **Restart** button.

Run Status: FINISHED
 Started: Jan 24, 2012 2:42:11 PM
 Current Step: Connectivity Test
 Options

Restart **Report**

Results Summary
 Finished Run + As of Jan 24, 2012 2:43:56 PM
 Page: 1 of 1

Step	Response	Message
Is latencyThreshold valid?	⊖	greater than
Is lossThreshold valid?	⊖	greater than
Operating System Detector	⊖	The OS system is Windows 2003.
Set Flow Variable	⊕	success
Ping	⊗	Connection Error
Set Context Value	⊕	success
Error : failure	⊗	Return step - Connectivity Test

History Tree
 Context
 Stack

Connectivity Test_1461
 Connectivity Test
 Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable.
 /Library/Accelerator Packs/Network/Connectivity/Connectivity Test

Graph
 Compass | Position | Open Graph | Tooltips

The graph shows a flow starting with 'Set Flow Variable' (success) leading to 'Ping' (failure). A 'diagnosed' step leads to another 'Set Flow Variable' (success) and 'Ping' (success), which then leads to 'Check L' (failure).

To restart the flow



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

Running subflows

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

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

To step into and out of a subflow

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




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



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

The **Step Out** button appears among the others.



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

OR

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

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

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

Previewing flows

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

The preview of the flow provides:

- A graph of the flow
- Buttons for starting the flow in various modes
- Various navigation icons.
- Panels with specialized information:
 - **Flow Details** panel: The flow's location in the Flow Library, its universally unique ID (UUID), and the flow's description, which includes a summary of what it does as well as important information for using the flow (such as the information required for inputs and the source and meaning of results).
 - **Reporting** panel: Charts built from data reported by flow inputs.
 - **Execution Links** panel: Links (URLs) for starting each kind of flow run. You can bookmark these links in your web browser or send them to another program or, in an email or a chat session, to another user. These links are URLs that launch the flow. If you change a run's name, its URLs in the execution links change accordingly.
 - **Graph** panel: Graphic depiction of the flow.

To preview a flow

- In the folder tree on the **Flow Library** tab, click the flow.

The flow preview is loaded and appears.

Preview Flow: Windows Health Check

Run name:

[Guided Run](#) [Run All](#) [Instant Run](#)

Flow Details

Windows Health Check
 /Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check

Flow UUID: 3541d63f-603a-449b-9d43-8e57d7d61482

Full Description:

This flow checks the overall health of a Windows server. It checks for many common problems, such as testing if the server is running, how much free memory the server has, and what percentage of its CPU is currently being utilized. A summary of the server's health is then either sent as email, displayed on the screen, written to a file or ignored.

NOTE:
 Only Windows XP, Windows 2003 and later versions track system uptime. Previous systems will always fail to obtain system uptime.

Inputs:
 host - The server to check the health of.
 user - The username to use when authenticating

Reporting

Execution Links

Graph







The graph shows a flow starting with 'Ping Target System' (green outline, success). It branches into 'Get System Uptime' (green outline, success) and 'Notify' (red outline, failure). 'Get System Uptime' leads to 'Get Host Name' (green outline, success). 'Notify' leads to 'Resolved' (green outline, success). 'Ping Target System' also leads to 'Create System Health Report' (green outline, success). 'Create System Health Report' leads to 'Get M...' (red outline, failure). 'Get M...' leads to 'Resolved' (green outline, success).

Tips:

- To return to the Library without running the flow that you are previewing, 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](#).

In the graph of the flow and the flow preview page, the symbols have the following meanings:

Symbol	Name	Meaning and comments
	Start step	The entry point for a flow. The green outline signifies that the step is the start step. (The flow operation symbol inside the step varies with the operation from which the step was made.)
	Parallel split step	A parallel split step is kind of hard to miss, because the lanes that follow tend to be rather large. Nonetheless, this symbol represents the start of the parallel split step. Parallel split steps are used to run distinct sequences of steps simultaneously.
	Multi-instance step	A multi-instance step performs the same action with different values for a given input (often specifying different targets), at the same time. For a multi-instance step, all the responses except one must ultimately loop back to the multi-instance step. The only response that leads ultimately to a return step is the group done response.

	Diagnosed return step	The step that ends a flow when a problem has been diagnosed.
	Error return step	The step that ends a flow when an error has occurred that prevents the flow from continuing. An error return step can also be used to indicate an expected condition is false (for instance if a counter is greater than a given number).
	Resolved return step	The step that ends a flow when a problem has been resolved.
	No action taken return step	The step that ends a flow when no action needs to be taken.
	Gated transition	A transition that is gated, or restricted to users with certain access permissions, appears in red on the canvas.
	Handoff transition	A transition after which the flow pauses for handoff to another user.

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

Navigating in the flow preview

To show more of the flow graph

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


To zoom in to 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.



The Zoom In, Restore, and Zoom Out icons

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.




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 positions of the panels on the preview page

- Click the **Reverse Columns** button.



To give the flow graph top-left orientation

- Click the **Position** button.



To see descriptions of each step

- Click the **Tooltips** button.



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

Deleting a run from within the run

Suppose you start a run and, for any number of reasons, you want to stop it. If you are in the flow preview, the best way to end the run is to click the **Delete** button.

To stop a run and delete its record

1. In the flow preview, locate the row of run-control buttons, and click <<more controls, to the right of the buttons.



More buttons appear.



2. Among the additional buttons, click the **Delete** button.

Updating a run


To update a run

1. In the flow preview, locate the row of run-control buttons, and click **more controls**, to the right of the buttons.



More buttons appear.



2. Among the additional buttons, click the **Update** button ().

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 browser window. The larger the flow, the more data that it generates, the more useful this can be. Consider the following example.

The screenshot displays the HP Operations Orchestration interface. At the top, there are navigation tabs: Dashboard, Flow Library, Current Runs, Reports, Scheduler, and Administration. The main area is divided into two panels. The left panel, titled 'Run Status: NEW', shows the current step 'Clear Context Value' and provides controls for 'Next Step', 'Run All', and 'Pause'. Below this is the 'Results Summary' section, which indicates 'Last Completed Step: No results yet' and includes a table with columns for 'Step', 'Response', and 'Message'. The right panel, titled 'Iterative Connectivity Test_901', shows a detailed flow graph. The graph consists of several nodes: 'Clear Context Value' (two instances), 'Error', 'Diagnosed', 'Iterator', and 'Connectivity Test'. The flow starts with 'Clear Context Value' leading to 'Error', which then leads to 'Diagnosed', and finally to 'Connectivity Test'. The 'Iterator' node is connected to 'Clear Context Value' and 'Error'. The 'Connectivity Test' node is connected to 'Diagnosed' and 'Error'. The graph shows various states: 'success', 'failure', 'diagnosed', 'has more', and 'no more'.

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

Central User Guide

Opening the flow graph in a separate browser window

The screenshot displays the HP Operations Orchestration Central interface. At the top, it shows 'Run Status: STOPPED' and 'Current Step: Clear Context Value'. Below this are control buttons for 'Next Step', 'Run All', and 'Pause'. A 'Results Summary' section indicates 'Last Completed Step: No results yet'. A table with columns 'Step' and 'Response' is visible, with a 'Display most recent: ON' toggle. The main area shows a flow graph with steps like 'Clear Context Value', 'Error', 'Diagnosed', and 'Iterator'. A separate browser window titled 'HP Operations Orchestration Central - Flow Graph View - Win...' is overlaid on the graph, showing a zoomed-in view of the flow graph with 'Size Window to Graph' and 'Dock Graph' buttons.

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

To open the flow graph in a different browser window

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

Seeing what has happened in a flow run

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

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

The screenshot displays the HP Operations Orchestration interface. On the left, the **Results Summary** panel shows the flow status as **STOPPED**, started on Jan 3, 2012 at 9:52:05 PM. The current step is **ConnectivityTest**. Below this, a table lists the steps:

Step	Response	Message
Clear Context Value	✓	success
Iterator	✓	Testing 16.59.62.5

The right side of the interface shows a **Graph** view of the flow. The graph includes nodes for **Iterator**, **Diagnosed**, **Error**, and **ConnectivityTest**. Transitions between these nodes are labeled with outcomes such as **success**, **failure**, and **diagnosed**. The **ConnectivityTest** node is highlighted with a red border, indicating it is the current step.

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

Notes:

- The OO group of which you are a member must have the **RUN_REPORTS** capability to view the **Reports** tab.
- To view runs other than your own on the **Current Runs** tab, your group must have the **MANAGE_RUNS** capability.
- The times given on the **Reports** and **Current Runs** tabs are given in the time zone for the Central client, whose relationship to Greenwich Mean Time (GMT) is shown in the Central web application's title bar.

To create a report of a flow run's history

Central User Guide

Seeing what has happened in a flow run

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

Select All	Name	Run ID	Type	Step	Started	Step Name	Step Start Time	Previous Step Name	State	User
<input type="checkbox"/>	Connectivity_Test_6	6	Interactive	0	Dec 30, 2011 3:42:31 PM	Is SQL Connected	Dec 30, 2011 3:46:1...	Connectivity Test	Stopped	admin
<input type="checkbox"/>	URL_Test_5	5	Interactive	9	Dec 30, 2011 12:28:04 PM	Resolved : success	Dec 30, 2011 1:24:3...	Display check	Input Required	admin
<input type="checkbox"/>	Windows_Health_Check_2	2	Interactive	3	Dec 30, 2011 12:14:07 PM	Resolved : success	Dec 30, 2011 12:22...	Evaluate Expression	Input Required	admin
<input type="checkbox"/>	Windows_Health_Check_1	1	Interactive	3	Dec 30, 2011 11:59:39 AM	Resolved : success	Dec 30, 2011 12:14...	Evaluate Expression	Input Required	admin

2. To see details about each step of the run, click the run name in the **Name** column.


The run's report on the **Reports** tab opens. The **Basic** tab on the run history report contains the number and name of each step that has completed up to this point along with their responses, messages, and call paths.

Step #	Step Name	Response	Message	Call Path
1	Is latencyThreshold valid?		greater than	Connectivity Test > Is latencyThreshold valid?
2	Is lossThreshold valid?		greater than	Connectivity Test > Is lossThreshold valid?
3	Operating System Detector		The OS system is Windows 2003.	Connectivity Test > Operating System Detector
4	Set Flow Variable		success	Connectivity Test > Set Flow Variable
5	Ping		Connection Error	Connectivity Test > Ping
6	Set Context Value		success	Connectivity Test > Set Context Value
7	Error : failure		Return Step - Connectivity Test	Connectivity Test > Error : failure

3. To see the details of each step that has been completed, click the **Advanced** tab.

Central User Guide

Seeing what has happened in a flow run



The screenshot shows a web interface for viewing a flow run instance. The title bar reads "Results > One Run Instance". Below this, a yellow box contains the following information:

- Run Id: 6
- Run Name: Connectivity Test_6
- Connectivity Test**
- Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable.
- [/Library/Accelerator Packs/Network/Connectivity/Connectivity Test](#)

Below the yellow box are two tabs: "Basic" and "Advanced", with "Advanced" selected. A "collapse all steps" link is visible above a table. The table has the following columns: Step Name, Step#, User Id, Parent Flow, Start Time, End Time, Execution Time [secs], Response, and Bound Inputs. The table contains one row for the "Connectivity Test" step.

Step Name	Step#	User Id	Parent Flow	Start Time	End Time	Execution Time [secs]	Response	Bound Inputs
Connectivity Test	0	admin		Jan 3, 2012 8:36:49 PM	Jan 3, 2012 8:51:33 PM	884.129		(latencyThreshold=4042), (lossThreshold=42), (password=*****), (host=16.59.62.4), (target=192.168.1.6), (username=cnc)

The **Advanced** tab may show more information than you need, but you can edit which types of information it shows, or hide it altogether.

4. To select which columns appear on the **Advanced** tab:
 - a. Scroll down the tab.
 - b. Select the names of the columns you want to appear under **Report Columns**.

The screenshot displays the 'Results > One Run Instance' window. At the top, it shows 'Run Id: 6' and 'Run Name: Connectivity Test_6'. Below this, a yellow box contains the title 'Connectivity Test' and a description: 'Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable. /Library/Accelerator Packs/Network/Connectivity/Connectivity Test'. The interface has two tabs: 'Basic' and 'Advanced', with 'Advanced' selected. Below the tabs is a table with one row of data:

Error : failure	7	admin	Connectivity Test	Jan 4, 2012 12:42:12 AM	Jan 4, 2012 12:42:12 AM
-----------------	---	-------	-------------------	-------------------------	-------------------------

Below the table is a 'Report Columns' section with a list of checkboxes for various report fields:

- Step Name
- Step#
- User Id
- Parent Flow
- Start Time
- End Time
- Execution Time [secs]
- Response
- Bound Inputs
- Recorded Bound Inputs
- Result
- Description
- ROI Value

At the bottom of the 'Report Columns' section, there is an 'Apply' button and a 'show all' link.

c. Click **Apply**.

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

For information on exporting a run history report, see [Exporting run history reports](#).

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

Run histories: What happened and why

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

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

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

- The time frame within which the flow or flows were run
- Central users who ran the flows, or OO authors who created or modified the flows
- Results of or actions taken by the flows
- The flow's location in the Flow Library
- Input name/value combinations
- Configuration items or domain terms that the flows reported to the Dashboard, including the flow's category

The screenshot shows the 'Reports' tab in the HP Operations Orchestration interface. The top navigation bar includes 'Dashboard', 'Flow Library', 'Current Runs', 'Reports' (selected), 'Scheduler', and 'Administration'. Below the navigation bar is a 'Filters' section with the following fields:

- Start Date:** 02/17/10
- End Date:** 03/02/10
- OR choose last...:** Radio buttons for 10 minutes, 12 hours, 7 days, 30 minutes, 24 hours, and 30 days.
- User:** Any user
- Results:** Any result
- Location:** Any location
- Inputs:** (empty)
- Domain Terms:** (empty)

Buttons for 'Update Results' and 'Reset' are located below the filters. Below the filters is a table titled 'Results > All Flow Types' with the following data:

Flow Name	Number of Runs	Avg. Number of Steps	Avg. Repair Time [secs]	Most Recent Run
Get CPU Performance	2	3	42149.978	02/18/10 10:46:50
Get Network Performance	1	3	25.686	02/18/10 10:31:18
Get Hard disk Performance	1	6	86.901	02/18/10 10:32:48
Service Status	1	3	39.499	02/18/10 10:41:13

At the bottom of the interface, there is a pagination bar showing 'Page 1 of 1', '50 Rows', and 'Export' options. The status 'Displaying 1 - 4 of 4' is also visible.

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

To define a search for run histories

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

You must specify a window of time for any search to produce results.

- To specify a fixed window of time:

- i. Click the calendar in **Start Date** and select a date.

OR

Type the date in the box, in the format mm/dd/yy.

- ii. To specify a time of day for the start of the window, press the TAB key on your keyboard and then, in the drop-down box to the right of the calendar, select a time from the list. Times are shown in 12-hour format, qualified by AM or PM.

- iii. Click the calendar in **End Date** and select a date.

OR

Type the date in the box, in the format mm/dd/yy.

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

OR

To pick a time window relative to the present, select an item from the list under **OR choose last**.

2. To specify runs or flows related to particular users or authors, in the **User** box:

- To specify runs or flows associated with any user, leave **Any User** in the box.
- To specify all the runs that you have started within the time window, select **My Runs**.

To specify only runs of a certain flow or flows, you can combine this filter with a location that you specify in the **Location** box.

- To specify all the flows whose content you have been the last person to modify, select **My flows**.
- To specify all the flows and runs executed by one or more other users:

- i. In the **User** box, select **Run by**.


- ii. In the **Run executed by users** box, type the usernames of users whose runs you want to look at, separating each two users by a comma (,).

3. To specify one or more particular results of the flow, in the **Results** drop-down list, select the result or results that you want to include in the list of flows that appears in the report.

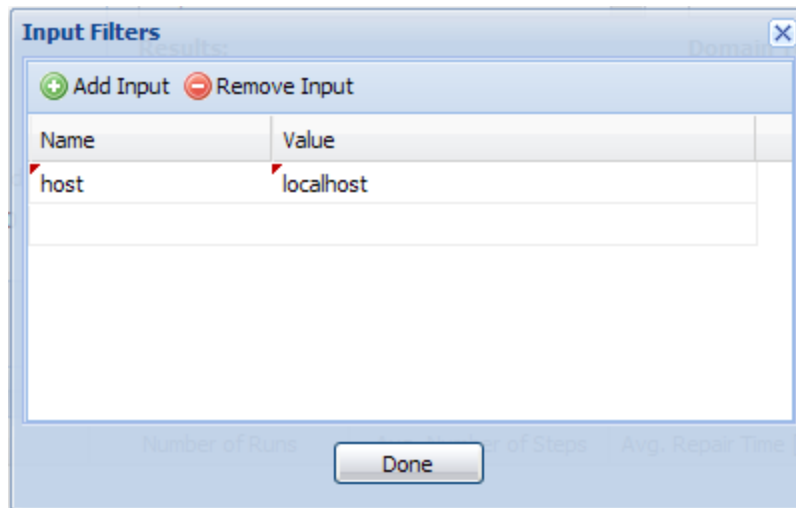
4. To specify the flows in a given folder, in the **Location** drop-down box, click the **Library** folder that contains the flows whose runs you want to see.

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

5. To specify flows that have flow or step inputs of certain values, in the **Inputs** box, add any input names and values you want to use to further limit the runs included in the report. You can specify only the input name. If you don't specify a value for the input, then the results will include all the flow runs of all the flows that have an input of that name, regardless of the value with which the flows ran.

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

The **Input Filters** box appears:



- i. To add an input, click **Add Input**.
- ii. Under **Name**, type the input name.
- iii. If you wish to specify a value for the input, then under **Value**, type the value.
- iv. Click **Done**.

OR

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

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

Where

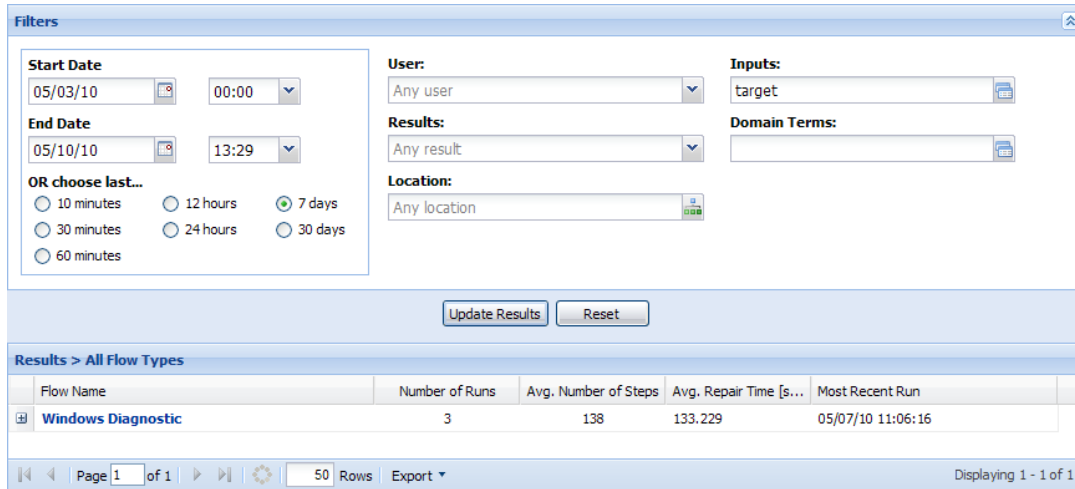
- <name> is the name of a flow or step input.
- <value> is the value for the input it is paired with. Note that specifying a value is optional.
- The separator between each <name:value> pair is a comma (no space).

To specify only an input name, without a value, the colon following the name is not needed.

6. To use one or more domain terms to further define the search, in the **Domain Terms** box, follow the same steps to specify domain term name/value pairs that you did for inputs.

7. Click **Update Results**.

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




The flows that have runs that fit the search criteria appear.

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

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

1. In the **Results** list, click the plus sign (+) next to the desired flow.

A pie chart that displays the results of the flow appears below the flow's name.



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

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

Exporting run history reports

On the **Reports** tab, you can at any time export the data in the report that is currently showing.

To export the data from the currently showing report

1. At the bottom of the **Reports** tab, click the **Export** button.

A list pops up, from which you can choose one of three formats to export the report as:

- **Text**
A formatted XML text file
- **Excel**
A Excel spreadsheet
- **Print**
A printer-friendly HTML file


2. Click one of the options.

Run histories in more detail

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

To view detailed information on flow runs

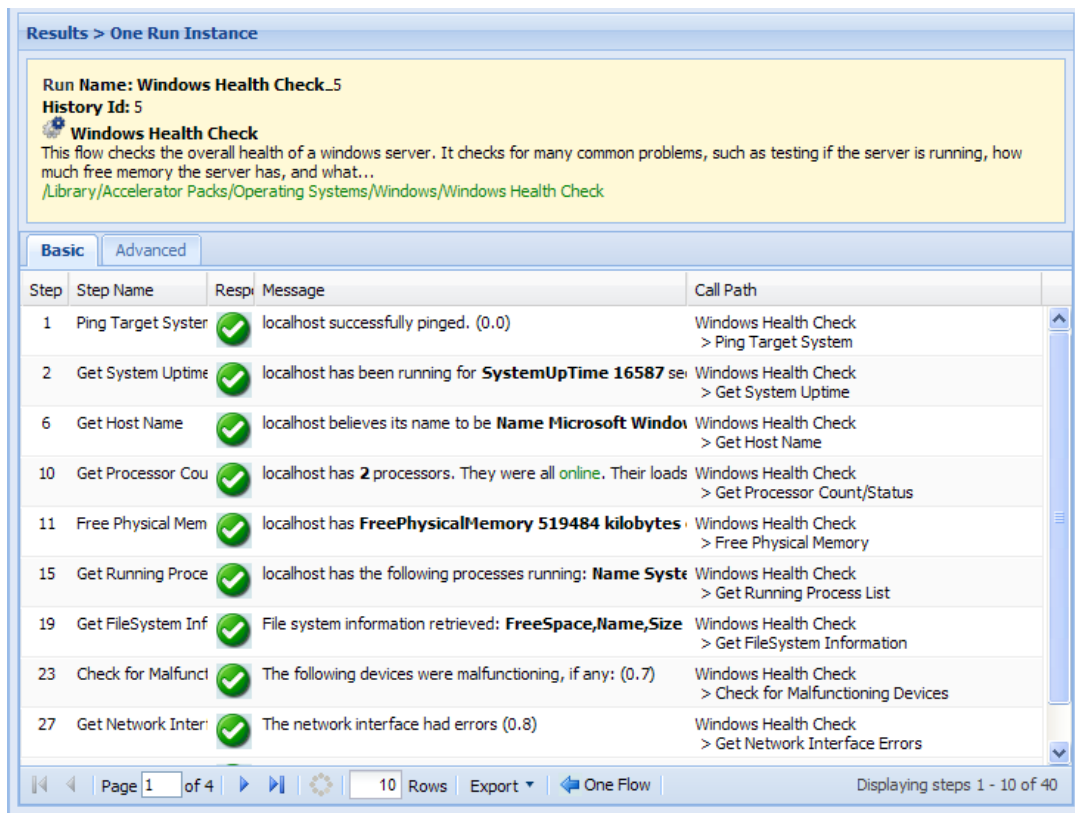
1. To view the list of runs for a given flow, click the flow name.



The screenshot shows a web interface with a header 'Results > One Flow Type'. Below the header is a yellow box containing the title 'Connectivity Test' and a description: 'Checks to see whether or not a connection exists, and if it exists checks to see if it is adequately reliable.' Below the description is a link: '/Library/Accelerator Packs/Network/Connectivity/Connectivity Test'. Below the yellow box is a table with the following columns: Run Id, Run Name, Start Time, Executed by, Time To Run, Result Type, Number o..., Last Modif..., Revision, and ROI Value. The table contains two rows of data.

Run Id	Run Name	Start Time	Executed by	Time To Run	Result Type	Number o...	Last Modif...	Revision	ROI Value
175	Connectivity Test_175	Dec 29, 2011 9:22:...	admin	74.452	Resolved	8	admin	81	0
134	Connectivity Test_134	Dec 29, 2011 9:10:...	admin	111.53	Resolved	8	admin	81	0


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



You can also specify which information is shown.

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

- Put the mouse on a column heading, and then click the downward-pointing arrow at the right of the column.
- Point to **Columns**, then in the list that appears, select the columns you want to see.



Viewing step results and other advanced data

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

To view advanced data for the steps of a flow run

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

The **Advanced** report appears:

The screenshot shows a web interface for viewing run history. At the top, it says 'Results > One Run Instance'. Below that, the run details are shown: 'Run Name: Windows Health Check...5', 'History Id: 5', and a description of the 'Windows Health Check' flow. The interface has two tabs: 'Basic' and 'Advanced', with 'Advanced' selected. Below the tabs, there is a 'collapse all steps' button. The main content is a table with the following columns: Step Name, Step#, Parent Flow, Start Time, and Response. The table contains three rows of data.

Step Name	Step#	Parent Flow	Start Time	Response
Windows Health Check	0		02/02/09 15:34:34	(password=(vitalProces (host=local
Ping Target System	1	Windows Health Check	02/02/09 15:35:02	success (localhost successfully pinged.) (host=local
Get System Uptime	2	Windows Health Check	02/02/09 15:35:07	success (localhost has been running for SystemUpTime 16587 seconds.) (format=cs Win32_Pert

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

2. To select which data are reported:
 - a. To collapse steps in the report, click **collapse all steps**.
 - b. Scroll to the bottom of the page.
 - c. Under **Report Columns**, select the boxes for the data that you want to see, unselect

those you don't, and then click **Apply**.

Basic **Advanced**

Report Columns

- Step Name
- Step#
- User Id
- Parent Flow
- Start Time
- End Time
- Execution Time [secs]
- Response
- Bound Inputs
- Recorded Bound Inputs
- Result
- Description
- ROI Value

[show all](#)

Scheduling flow runs

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

There are a couple of requirements for scheduling a flow:

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

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

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

You can create and manage flow schedules by using the WSCentral Service SOAP API. For information on using the SOAP API, see the *OO Software Development Kit Guide* (SDKGuide.pdf). To use the WSCentralService SOAP API to access Central features, you must be logged in with an account that has the `HEADLESS_FLOW` capability. For information on capabilities, see the *OO Administrator’s Guide* (AdminGuide.pdf).

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.

Note: 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.

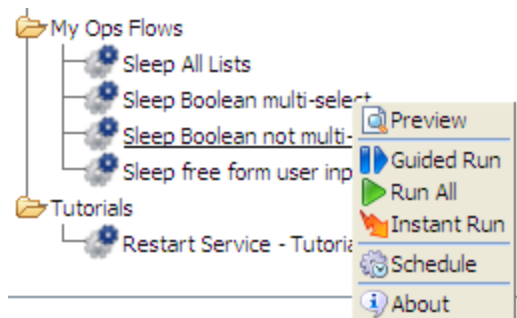
For information on disabling run concurrency see the *OO Administrator Guide* (OOAdminGuide.pdf).

Creating a schedule for a flow

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

To create a schedule for a flow

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



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

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

Schedule Flow

Schedule | Inputs

Flow

Flow Name: /Library/Accelerator Packs/Network/Connectivity/Connectivity Test

Schedule Name: Connectivity Test

Recurrence pattern

Daily Every 1 minutes

Weekly Every day

Monthly Every weekday

Yearly

Range of recurrence

Start date: 01 10, 2012 No end date

Start time: 12:00 AM End after 1 occurrences

End by 01 10, 2012

Save Schedule Cancel

2. In the **Schedule Name** box, type a name for the schedule, or accept the default name, which is the name of the flow.

If you leave the box blank, the schedule name is the default name. The schedule name is given to all the runs in the schedule.

The schedule's name is given to all the runs that the schedule triggers. If you include the variable {runID} in the schedule name, then each of the schedule's runs is distinguished by its runID.

Note: The schedule name cannot include a colon (:).


3. On the **Schedule** tab, define the flow's schedule for running: the time(s), frequency, and starting and stopping points of the schedule.

- a. Under **Recurrence Pattern**, specify how often you want the flow to run.

The **Recurrence Pattern** area changes to present options with which you can specify **Daily**, **Weekly**, **Monthly**, or **Yearly**, frequencies of recurrence.

- b. Under **Range of Recurrence**, specify the schedule's starting date and time and specify whether and how its end is determined.

4. To specify values for the flow-level inputs, click the **Inputs** tab.




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

5. To specify a value for an input, double-click the input's **Value** box, and then do one of the following, according to whether the value is typed or chosen from a list:

- Type the value.

if the value must be encrypted (most commonly, the password input for a set of credentials), a padlock icon (🔒) appears in the **Encrypted** column of the input's row. The **Input Encrypted Value** dialog appears, and you type the value in this dialog, as in the following example:



Type the value in the text box(es), and then click **OK**.

- If the **List Input** dialog appears when you put your cursor in an input's box, then in the **List Input** dialog, click the downward-pointing arrow, select the desired value from the drop-down list, and then click **OK**.
6. 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 you see specified.

The separator character is probably specified as a **delimiter** or **separator** input. 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**


Be sure not to include a space between the values unless the defined separator character is a sequence that includes a space.

You can remove inputs that are not required.

7. To remove an input, click the red X (✗) in the input's row under **Remove**.
8. To add an input:
- On the **Inputs** 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.

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

The following appears.



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

Note: If you are scheduling a flow that contains a Notify or Basic Notify step, you must enter a value for the notifyMethod input (for instance, **display** or **file**) as well as for any other required inputs. This is true for flow inputs as well as step inputs. For example, if you give the notifyMethod input a value of **file**, then you must also give the destination input a value.

If the Notify or Basic Notify step inputs are step inputs rather than flow inputs, you add the inputs in the scheduler and specify values for them as described in step 7 in the preceding procedure.

Editing flow schedules

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

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

To edit a flow's schedule

1. On the **Scheduler** tab, click the flow whose schedule you want to change.

The screenshot shows the Scheduler tab interface. At the top, there are navigation tabs: Dashboard, Flow Library, Current Runs, Reports, Scheduler (selected), and Administration. Below the tabs are buttons for Refresh and Disable All. The main area displays a table titled "All Flows With Schedules".

Scheduled Flow	Enabled	Controls	Previous run time	Next run time	Delete
Windows Health Check /Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check				Jan 4, 2012 11:30:00 PM - "Windows Health Check"	
Iterative Connectivity Test /Library/Accelerator Packs/Network/Connectivity/Iterative Connectivity Test			Jan 4, 2012 11:15:00 PM - "Iterative Connectivity Test"	Jan 4, 2012 11:30:00 PM - "Iterative Connectivity Test"	

Below the table, there is a note: "Click on a Flow name to see all the schedules running for that flow."

2. To see the schedule(s) for a single flow, click the flow name.

The screenshot shows the Scheduler page for a specific flow: "Windows Health Check". The page title is "Schedules For One Ops Flow". It includes a "Scheduler Page" link and a "Refresh" button. Below is a table with columns: Edit Schedule, Schedule Name, Starting Date/Time, Ending Date/Time, Recurrence, Enabled, Controls, Previous run time, Next run time, Parameters, and Delete.

Edit Schedule	Schedule Name	Starting Date/Time	Ending Date/Time	Recurrence	Enabled	Controls	Previous run time	Next run time	Parameters	Delete
	Windows Health Check weekly	Jan 5, 2012 12:00:00 AM	Jan 16, 2012 12:00:00 AM	every Wed, Thur				Jan 11, 2012 12:00:00 AM	host=16.59.62.4 notifyMethod=Display password=***** packetSize=420 packetCount=420 ...	
	Windows Health Check	Jan 4, 2012 12:00:00 AM	Jan 6, 2012 12:00:00 AM	every 15 minutes			Jan 5, 2012 1:00:00 AM	Jan 5, 2012 1:15:00 AM	host=16.59.62.4 notifyMethod=Display packetSize=420 password=***** packetCount=420 ...	

Note: The inputs that were specified for each schedule are shown in the **Parameters** column. To see more of the inputs, widen the column.

3. To make changes to the schedule, click the clock icon () in the **Edit Schedule** column.
4. In the **Schedule Flow** dialog, 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](#).

If you change the schedule name, the new schedule name is given to all flow runs in the schedule from this point on.

Enabling and disabling flow schedules

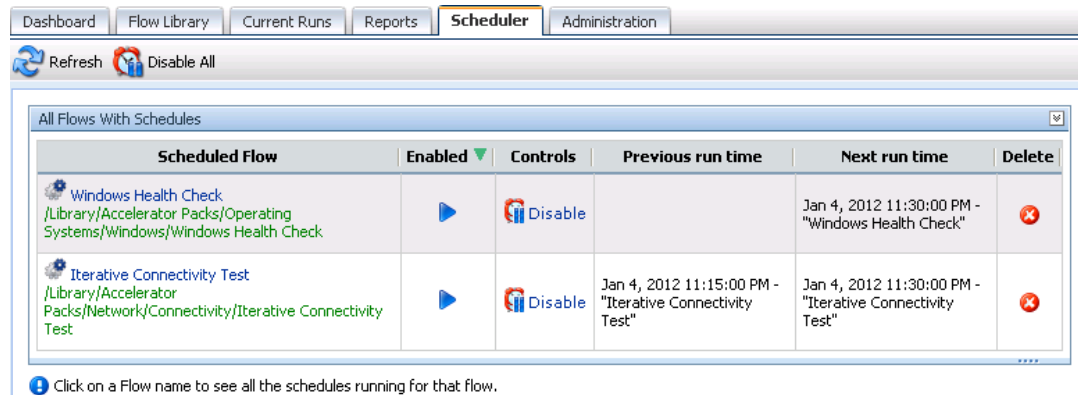
The main page of the **Scheduler** tab lists the flows that have schedules created for them.

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

To enable or disable run schedules for a flow

1. Click the **Scheduler** tab.

The flows that have one or more schedules are listed, with information and controls displayed.



2. To enable or disable all of a flow's schedules, click **Enable** or **Disable** in the **Controls** column, as appropriate.

OR

To enable or disable individual schedules for a flow that has more than one schedule:

- a. To see the schedule(s) for a single flow, click the flow name.

The schedules for the flow appear.

Edit Schedule	Schedule Name	Starting Date/Time	Ending Date/Time	Recurrence	Enabled	Controls	Previous run time	Next run time	Parameters	Delete
	Windows Health Check weekly	Jan 5, 2012 12:00:00 AM	Jan 16, 2012 12:00:00 AM	every Wed, Thur				Jan 11, 2012 12:00:00 AM	host=16.59.62.4 notifyMethod=Display password=***** packetSize=420 packetCount=420 ...	
	Windows Health Check	Jan 4, 2012 12:00:00 AM	Jan 6, 2012 12:00:00 AM	every 15 minutes			Jan 5, 2012 1:00:00 AM	Jan 5, 2012 1:15:00 AM	host=16.59.62.4 notifyMethod=Display password=***** packetSize=420 packetCount=420 ...	

To return to the display of all the flows, click **Scheduler Page** (Scheduler Page).

- b. In the row of the schedule you want to delete, click **Enable** or **Disable** in the **Controls** column.

Deleting flow schedules

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

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

To delete run schedules for a flow

1. Click the **Scheduler** tab.

Scheduled Flow	Enabled	Controls	Previous run time	Next run time	Delete
Windows Health Check /Library/Accelerator Packs/Operating Systems/Windows/Windows Health Check				Jan 4, 2012 11:30:00 PM - "Windows Health Check"	
Iterative Connectivity Test /Library/Accelerator Packs/Network/Connectivity/Iterative Connectivity Test			Jan 4, 2012 11:15:00 PM - "Iterative Connectivity Test"	Jan 4, 2012 11:30:00 PM - "Iterative Connectivity Test"	

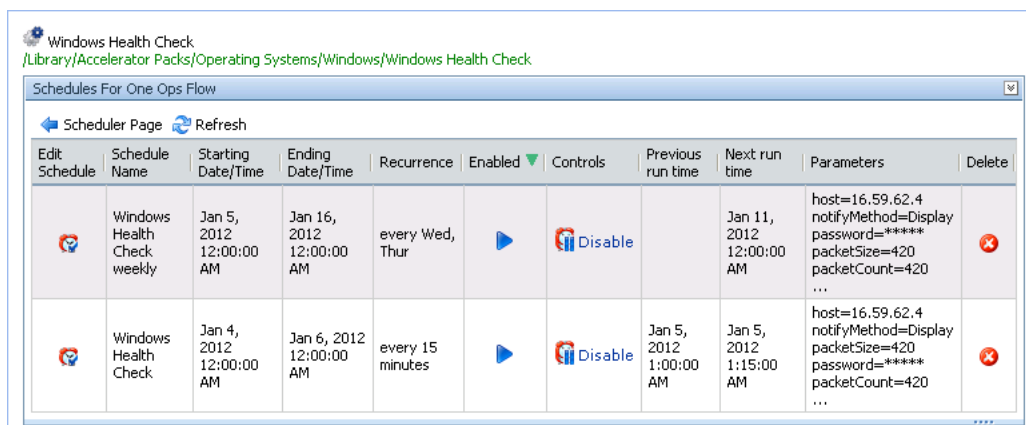
Click on a Flow name to see all the schedules running for that flow.

2. To delete all of a flow's schedules, click the red ball () in the **Delete** column for the flow.

OR

To delete individual schedules for a flow that has more than one schedule:

- a. To see the schedule(s) for a single flow, click the flow name.



To return to the display of all the flows, click **Scheduler Page** (Scheduler Page).

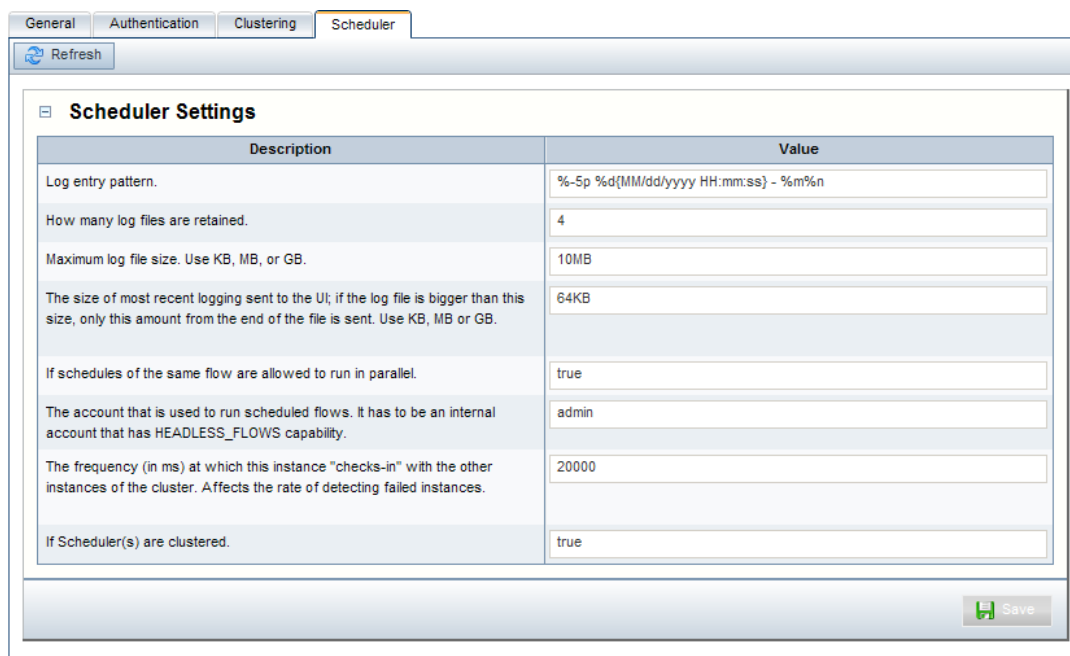
- b. In the row of the schedule you want to delete, click the red ball () in the **Delete** column.

Configuring Scheduler settings

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

To specify Scheduler settings

1. On the **Administration** tab, click the **System Configuration** subtab, and then click the **Scheduler** tab.



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

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	The number of logs that are stored on the system.	This setting affects the amount of space that is needed for log data storage.
Maximum log file size	The amount of space needed for log data storage.	Specify the maximum log file size in KB, MB, or GB.
The size of the most recent logging sent to the UI...	Limits how much log data is presented in the Scheduler UI.	Specify in KB, MB, or GB.
If schedules of the same flow are allowed to run in parallel.		The options are "true" or "false". The default is "true".
The account that is used to run scheduled flows.		By default, this is the "admin" account.
The frequency (in ms) at which this instance "checks-in" with the other instances of the cluster...		Because the Scheduler is entirely part of Central, this is redundant with the clustering configuration. Take no action.
If Scheduler(s) are clustered		Because the Scheduler is entirely part of Central, this is unnecessary. Always leave set to "true".

The Dashboard: Learning more from flows

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

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

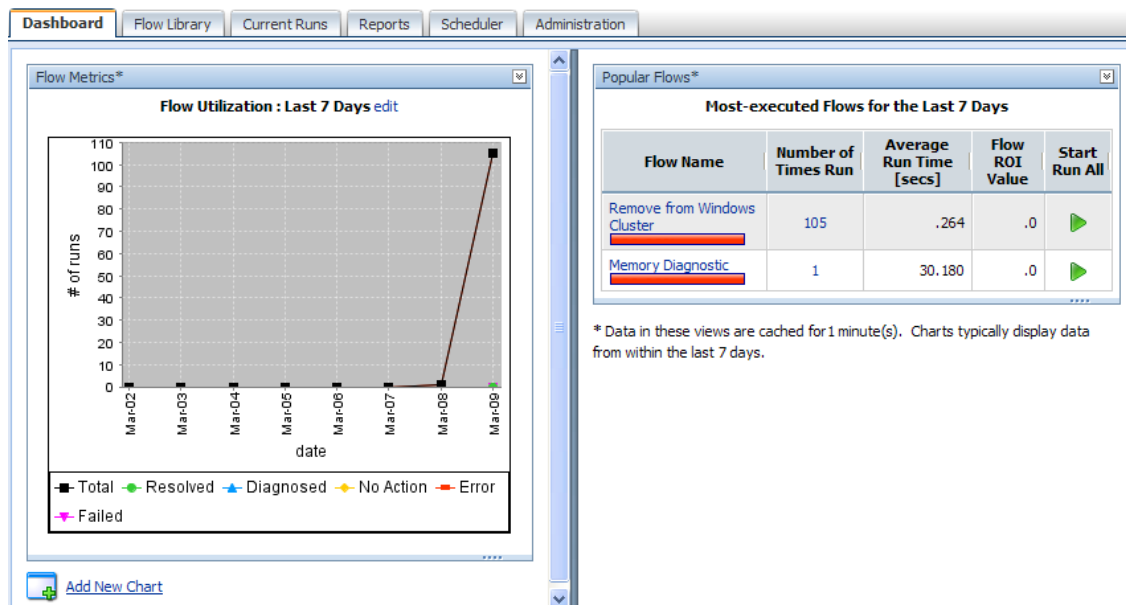
Answer: With Dashboard reporting charts and run-history reports on the **Reports** tab.

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

Collating data on Dashboard reporting charts

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

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



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

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

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

- Which server or application generated the most alerts?

The **Alerts per Configuration Item** chart answers both those questions.

- What actions have been taken on each application and server?

Look at the **All CI's organized by Action** chart.

- Which flows have run the most often, and what were their outcomes?

Consider the **Outcomes per Flow** chart.

- Which flows were run to resolve errors, and how many times did the flows run?


Open the **All Alerts of Severity=Error Resolved by Flows** chart.

Most of these charts require that the data they use are reported from the flow runs. Data from flow runs is reported from 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.



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

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

- Configuration Item (aka CI)

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

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

- Your company home page:
 - **CI**: the home page’s URL
 - **CI Type**: “Application”
 - **CI Minor Type**: “Web Page”
- Categories


The groups to which flow authors assign flows. Charting categories enables you to view performance of these classes of flows. See Studio Help for more information.
- Alerts, Incidents, Problems
 - Alerts are monitoring messages about possible error states that have arisen amidst IT operations.
 - Incidents can represent trouble tickets in Incident Management or trouble-ticketing systems that you run.
 - Problems can represent items in any Problem Management system you operate.
- Actions

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

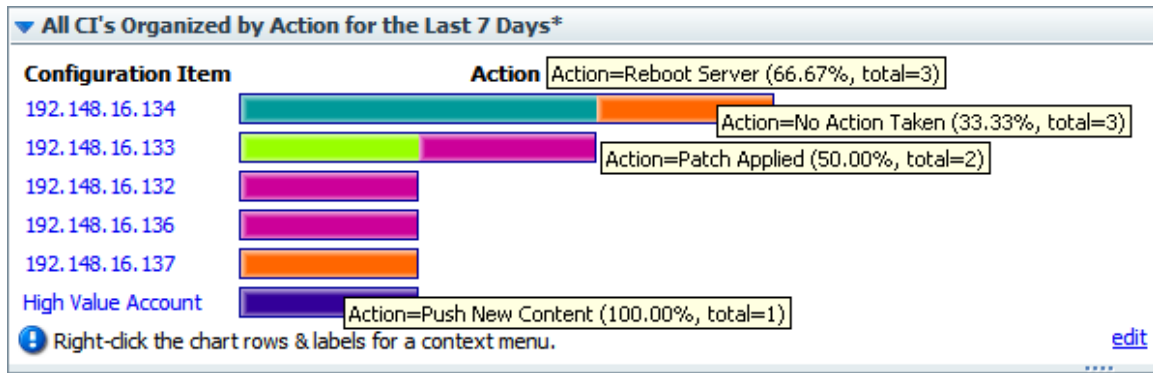
Outcomes are the return states of flows: Resolved, Diagnosed, No Action, Failure

What do the bars tell you?

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



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.



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

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

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

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

We can go further by drilling down into individual bars. For how we can learn more, see the following topic, [Learning more from the charts](#).

Learning more from the charts

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

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



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


To learn more about data items in a chart

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

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

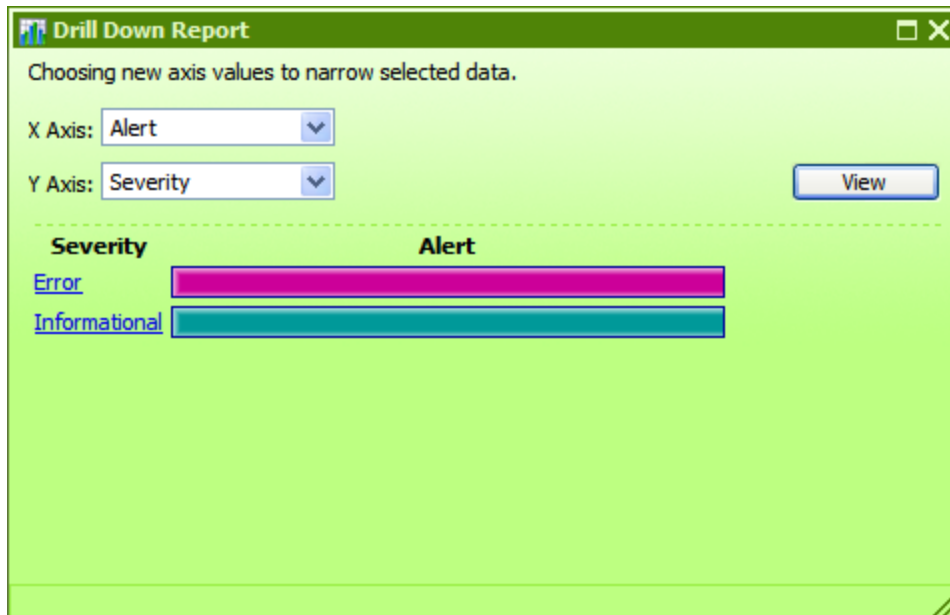
For instance, to show the distribution of alerts and severity levels for the Reboot Server action, right-click the bar segment representing Reboot Server, and then click **Drill Down**.

The following box appears.



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.



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. For more information on creating your own charts, see the following topic, [Creating and modifying charts](#).

Creating and modifying charts

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

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

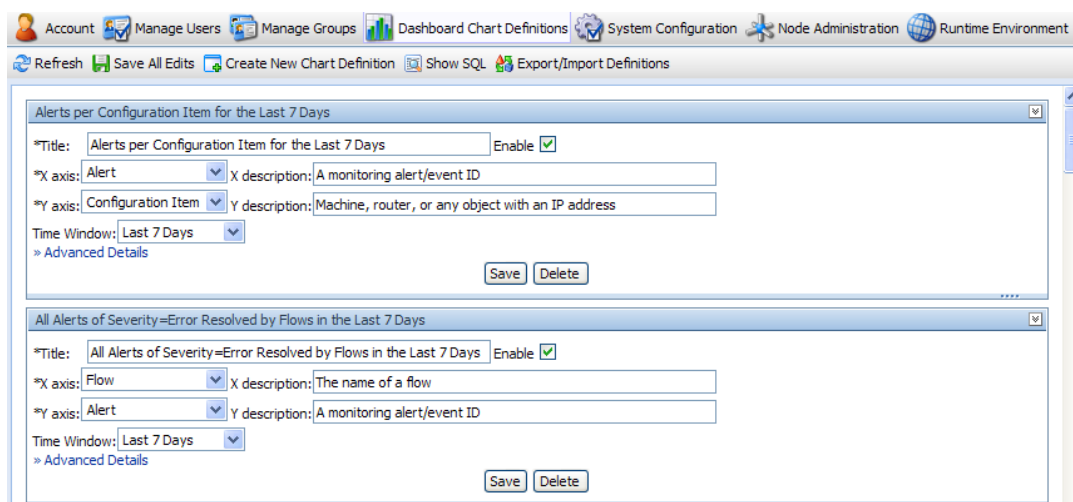
You redefine existing charts or create new charts on the **Administration** tab, by specifying which information is charted on the horizontal (X) and vertical (Y) axes.

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

To create a Dashboard chart

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

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



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

A new chart definition box appears.

The screenshot shows a form titled "Add New Item Here". It contains the following fields and controls:

- *Title: [text input] Enable
- *X axis: [Action] X description: [text input]
- *Y axis: [Action] Y description: [text input]
- Time Window: [Last 24 Hours]
- »Advanced Details (link)
- Save (button) Delete (button)

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**.

The screenshot shows the "Add New Item Here" form with the "Advanced Details" section expanded. It includes the following fields and controls:

- *Title: [text input] Enable
- *X axis: [Action] X description: [text input]
- *Y axis: [Action] Y description: [text input]
- Time Window: [Last 24 Hours]
- »Advanced Details (link)
- Comments: [text input]
- Top X: [0] X Threshold: [0] Top Y: [0] Auto-group
- Restrict X axis values: [text input]
- Group all other values in one segment:
- Name for all other values segment: [Others]
- Additional Constraints...**
- Domain Term Name: [text input] Domain Term Value: [text input]
- Save (button) Delete (button)

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

8. To chart only the most common occurrences of the element you're charting on the X axis or Y axis, type a number in **Top X** or **Top Y**.

For instance, to have this chart show you only the three most common types of alerts, type **3** in the **Top X** box.

9. To establish a floor value below which the X axis element is not charted, type the floor value in

X Threshold.

For example, to leave uncharted any alert types that don't have at least five instances reported, type **5** in the **X Threshold** box.

10. To chart only elements of a certain type (as represented by a domain term value), type the domain term value in the **Restrict X axis values** box.

So suppose you want to chart only alerts of the Loss of Connectivity type. Assuming that the flow author has created a domain term for Loss of Connectivity, you could type **Loss of Connectivity** in the **Restrict X axis values** box.

Besides restricting the Y axis to the most common occurrences of the element charted there, you can further restrict what is charted in the Y axis.

11. To group any values outside those you have specified into a single color segment of the bar on the X axis, select the **Group all other values in one segment** box.
12. If you have grouped into a single X-axis segment values outside those you have specified, name the segment by typing a name in the **Name for all other values segment** box.
13. Under **Additional Constraints**, from the **Domain Term Name** drop-down, select the domain term charted on the Y axis, and then type a value in the **Domain Term Value** box.
14. In the bottom of the chart definition box, click **Save**.

Let's look at our examples:

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

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

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

To edit a Dashboard chart definition

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

Notes:

- Using these charts requires that the flows reported have their relevant inputs configured to report data in the domain terms that the charts need. For information on how to add this reporting capacity to inputs, see Help for Studio.

- You can add new terms that you want to appear in the **X Axis** and **Y Axis** drop-down lists. To learn more about this see the material about domain terms in Help for Studio.

Exporting and importing chart definitions

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

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

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

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

To export chart definitions

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



3. In this dialog, click **Export**.
The **File Download** dialog box appears, for the XML file that stores the chart definitions.
4. In the **File Download** dialog, click **Save**.
5. In the **Save As** dialog box that appears, navigate to the location where you want to store the XML file, and then, if you want to rename it, specify another name.
6. In the **Download complete** dialog box that appears to tell you that the export has finished, click **Close**.

OR

Open the file to examine the definitions.

To import chart definitions

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



3. In the **Import** text box, type a path and filename for the chart-definition XML file that you want to import.
OR
Click **Browse**, and in the **Choose File** dialog, navigate to the XML file of definitions and click **Open**.
4. Back in the **Export/Import Definitions** dialog, click **Load File**.

Starting a flow from outside Central

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

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

These techniques are described in the *OO Software Development Kit Guide (SDKGuide.pdf)*.


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

Creating a link to a flow run

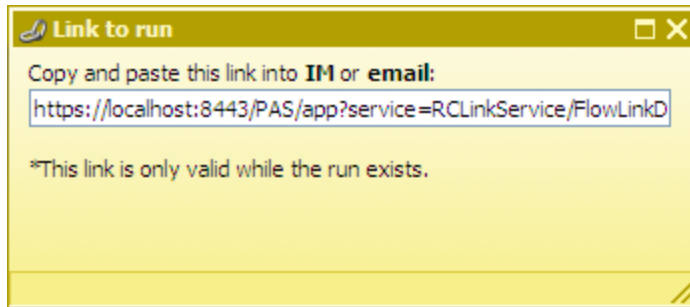
You can also enable another Central user to resume a flow run that you have paused (or resume the run yourself) by capturing the run's URL mid-run, then placing the URL on another Web page or in an IM or e-mail message. The URL that you capture is not the same as the address that is in your Web browser address box during the run.

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

To create a link to a flow run

1. When previewing a flow or during a run (except an instant run), above the run-type buttons in the upper-left panel of the flow run page, click the Options button (.
2. In the menu that appears, click **Link to this run**.

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




3. In the **Copy and paste this link...** text box, copy the entire URL, and then close the **Link to run** dialog box.
4. In the external source in which you are creating access to the run, paste the shortcut that you copied.

Creating a link to a flow

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

To create a link to a flow

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



2. At the bottom of the pane, click the **Execution Links** bar.
3. In the **Execution Links** section that appears, links that you can paste into a Web browser's

URL address line to start the flow appear, labeled for the kind of run you can carry out with them.

Execution Links

The below hyperlinks provide a valid URL for bookmarking, sending in email, etc. Use these links (by right-clicking to obtain your browsers' context menu) instead of the address in the browser address box.

Guided Run (step-by-step mode)
`https://localhost:8443/PAS/app?service=RCLinkService/FlowLinkDispatch&sp=SNEWRUN&sp=S354`

Run All (run-to-completion mode)
`https://localhost:8443/PAS/app?service=RCLinkService/FlowLinkDispatch&sp=SNEWRUNALL&sp=S`

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

Handing off flow runs

You might need to hand off a flow in which:

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

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


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

To hand off a flow run

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

The run is paused, and the state of the run is changed to **Handed Off**. Focus in Central shifts to the **Current Runs** tab.

A new e-mail message appears, with the URL of the flow included in the body of the message. The run name appears in the subject line.



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

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

1. Open the e-mail message that contains the URL of the flow and click the URL.

A new browser instance of Central opens, with the flow run ready to resume. You can get it going again as you would any other flow run that is paused.

2. To continue the flow run, click either the **Next Step** or the **Run All** button.

Auditing and managing flows

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

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

Users, groups, and access control

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

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

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

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

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

For more information, see the *OO Administrator Guide* (OOAdminGuide.pdf).

Capabilities and access permissions

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

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

that they use to author flows. For instance:

- To debug a flow, an author must have the Execute permission for that flow.
- A flow author must have the Link to permission for any flow or operation from which he or she creates a step in a flow.
- To change a system account, an author must have the Read and Write permissions for the system account.

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

Capabilities

Following are the capabilities that can be assigned in OO.

Capability	Description
MANAGE_USERS	Create, delete, and modify internal users and map external user groups (groups that exist outside of OO) to internal OO groups. Only holders of this capability can create internal OO users.
MANAGE_GROUPS	Create, delete, and modify groups.
AUTHOR	Start Studio.
SCHEDULE	Schedule flows.
MANAGE_RUNS	View, start, stop, delete, and reassign runs other than the user's.
RUN_REPORTS	View reports and Dashboard metrics/charts.
MANAGE_CONF	Manage configuration properties and dashboards.
VIEW_SCHEDULES	View flow schedules.
HEADLESS_FLOWS	Start flows from outside Central.

Permissions

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

Permissions for OO objects

Central User Guide

Users, groups, and access control

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

OO objects and the permissions needed to work with them

Object	Action	Necessary permission(s)
Folder		
	View contents	Read
	Add to contents	Read, Write (also needed for all children of the folder)
	Move	Read, Write
	Rename	Read, Write
Flow or operation		
	View/open	Read
	Modify	Read, Write
	Rename	Read, Write
	Execute/Run	Read, Execute
	Use as a step or subflow	Read, Link to
System account		
	View account name	Read
	Change account password	Read, Write
	Rename account	Read, Write
	Use in flow or operation	Read, Link to
	Use at runtime	Read, Execute

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

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

Allowing external users into the Central system

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

Logging on to OO has two phases:

- **Authentication:** This phase requires that the user's credentials (username/password) are verified against a provider, either OO (in the case of internal OO accounts) or an external provider.
- **Authorization:** This phase consists of determining an authenticated user's group memberships. Group membership is the basis of security in OO as permissions (access control lists) and capabilities are assigned to groups and NOT to individual users. A user cannot log on if his group set is empty.

Notes:

- Any user that is authorized to use OO by passing the above phases, is automatically assigned to the EVERYBODY group.
- It is unlikely that you will want all authenticated users in a large organization to use Central or Studio. So, when using an external authentication provider, you can restrict the set of users authorized to use OO by one of the following methods:
 - Configuring the provider so that only a subset of users are considered valid OO users. For example, in the case of Microsoft AD or LDAP, you can consider valid OO users those who are located in particular directory contexts.
 - Explicitly denying access to OO to a particular external user by defining an internal OO account that EXACTLY matches the fully qualified username—"DOMAIN\username" for AD, "username@realm" for Kerberos, or just "username" for LDAP—and then marking that account as disabled.

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

Using external authentication for Central users

To authenticate with any or all of Active Directory (AD), LDAP, or Kerberos authentication providers, you use the **Administration** tab's **System Configuration** subtab. The **System Configuration** page contains an **Authentication** tab that includes a section for each of the three kinds of authentication providers.

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

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

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

The information in these sections may look a bit formidable if you are not an AD, LDAP, or Kerberos administrator, but we will use the data in the following example throughout the sections to 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.

Example data

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

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

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

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

AD authentication settings (Deprecated)

Note: The AD authentication settings are deprecated, beginning with in OO 9.0. You should use the [LDAP authentication settings](#).

The following procedure for authenticating with AD refers to the **Authentication** tab on the **System Configuration** page.

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

To authenticate with Active Directory

1. Click the **Administration** tab, and then click the **System Configuration** tab.
2. On the **System Configuration** page, click the **Authentication** tab.
3. Scroll to the **AD Settings** area.

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 "domainusername").	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

4. Select the **AD Enabled** check box.
5. In the **Value** box of the default group row, specify the group to which unmapped OO groups or individuals are mapped to when mapping is not specified.
For instance, to assign unmapped groups or individuals only the capability to run flows, you would type `LEVEL_ONE`. To learn about the default OO groups, see [Managing groups](#).
6. 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"

So Tom 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:

Central User Guide

Users, groups, and access control

```
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.
7. In the **Value** box of the **List of LDAP contexts containing user groups** row, 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 semicolons (;) to separate contexts.

For example, suppose that:

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

The following setting makes these groups visible to OO:

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

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

8. In the **Value** box of the **Attribute of any group** row, leave the value **name**.
9. In the **Value** box of the **Active Directory URL** row, enter the URL or IP address of the AD, using the following syntax:

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

where:

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

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

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

OR

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

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

Central User Guide

Users, groups, and access control

(OOAdminGuide.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 value would be:

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

OR

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

10. In the **Value** box of the user domain row, type the domain where the users reside using the following syntax:

```
<domain>\{0}
```

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

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

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

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

Using this procedure's example, the **AD Settings** area should look like the following.


The screenshot shows the 'AD Settings (Deprecated, use LDAP)' configuration page. The 'AD Enabled' checkbox is checked. The settings are organized into a table with two columns: 'Description' and 'Value'.

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

At the bottom of the form, there is a status bar with a warning icon and the text 'There are unsaved changes on this Form', along with 'Test' and 'Save' buttons.

12. To save your settings, click **Save**.
13. To test the current AD authentication settings from within OO, click the **Test** button.

The **Test AD Settings** dialog box appears.



Test AD Settings

Enter a sample user to Test Authentication with your new settings:
*These credentials are not stored

Username:

Password:

Test Close

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

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

LDAP authentication settings

Note: LDAP providers should be used for authenticating against Microsoft Active Directory (AD) as well as the AD Global Catalog, and Microsoft Active Directory Application Mode (ADAM).

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

LDAP authentication requires the user who is logging on to have search capabilities or to use an LDAP manager account. If the user logs on with an LDAP manager account, the account must be defined as an internal OO system account (preferably a disabled account), and that the username must be a full LDAP Distinguished Name (DN). The user does not need to log on with an LDAP manager account if he has search capabilities.

If a user logs on using an attribute other than the DN (this is very common for non-AD LDAP directories), he should use an LDAP manager account because an LDAP directory can only be bound by using the user's DN.

For example, if the user's DN is:

```
CN=Tom Gage,CN=Users,DC=mirage,DC=com
```

and he uses the attribute `UID=tgage` to log on, the system must first look up Tom's DN, and then attempt to bind the directory with his DN and password. This requires an LDAP manager account even if Tom has search capabilities.

LDAP authorization consists of collecting the user's groups. If a manager account is set up, that account is used, otherwise the context of the user logging on is used for group lookup.

You can configure multiple chained LDAP providers (two or more directory servers chained together using LDAP, where LDAP requests are passed from one directory server to another) by clicking the **Add LDAP Configuration** button. Each configuration follows the same rules. In the case of non-AD LDAP directories (which do not have the concept of a domain), multiple LDAP providers can be chained to allow users from multiple directories to log into OO. In this case, there should be no name collisions between users across the different directories.

When an LDAP configuration's target is a Microsoft AD, an optional extra field, AD domain, is enabled. You can then enter the domain name in that field. If there are multiple LDAP domain configurations defined, or if there is only one AD and the AD Domain field is not configured, you must use DOMAIN\username to log on. You can use multiple AD domains by any of the following methods:

- Using multiple chained providers.
- Using an AD Global Catalog.
- Using a combination of multiple chained providers and an AD Global Catalog.

Every time there is more than one domain involved, the user must provide the FQN DOMAIN\username when logging on.

If there is only one AD configuration and the AD Domain field is configured, then the user can log on just by typing his username.

In OO 9.00, users are now tracked by the FQN:

- If the user has an internal OO account, then FQN = username.
- If the user is an LDAP user (non-AD LDAP), then FQN = username.
- If the user is an AD user, then FQN = DOMAIN\username.
- If the user is a Kerberos user, then FQN = username@realm.

The FQN will appear in all aspects of OO, such as OO Central, Studio, and repository workspaces.

This means that if an AD or Kerberos user checks out work before upgrading to OO 9.0, and then upgrades to OO 9.0, the user workspace is in the new FQN notation. For the user to be able to check in his work after the upgrade, he has to change his user folder name under to use the new FQN notation.

For instance, if a user logged on to an earlier version of OO as the Kerberos user tgage against the domain ssorealm, his user folder was called tgage; after upgrading to OO 9.0, the folder name becomes ssorealm%5Ctgage (the folder name uses "%5C" instead of "\" because a folder name cannot contain a backslash). The user must then rename his user folder in the file system to the match the new name format.

Caution: Because of this fact, it is very important for users to check in all of their work before upgrading to OO 9.0.

Regular LDAP configurations and AD user accounts that have restricted rights and cannot look up their own groups (for example, those that have no LDAP list privileges), should use LDAP manager accounts.

You can use an OO account (marked as external) to override an LDAP/AD user. The override can be used to:

Central User Guide

Users, groups, and access control

- Explicitly deny OO access to a user that otherwise would be allowed to log on by using the external provider(s) and the group mapping rules.
- Give a particular user a group membership set regardless of the group mappings defined in the system.

In OO 9.0, users are tracked by FQN, so the overriding account must use FQN. This means that if in an earlier version of OO, the AD user "tgage" in the MIRAGE domain could be overridden by defining an external OO account named "tgage". In OO 9.0, the overriding account has to be "MIRAGE\tgage". This applies whether or not the user has to use FQN when logging on to OO in the case of a single AD domain.

The LDAP provider marked as AD can be used against an AD Global Catalog (GC); in this case you should leave the **Value** box in the **AD Domain** row empty. For example, you can set the URL to point to `ldap://kdc.domain.com:3268` (domain forest Global Catalogue). Fully Qualified Names must be used for logging on any time more than one domain is involved.

The LDAP provider marked as AD can also be used to work with Microsoft ADAM, while the deprecated AD provider cannot be used with ADAM.

To authenticate with LDAP

1. Click the **Administration** tab, and then click the **System Configuration** tab.
2. On the **System Configuration** page, click the **Authentication** tab.
3. Scroll down to the **LDAP Settings** area.

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

4. Select the **LDAP Enabled** check box.
5. If the target is a Microsoft AD, select the **Microsoft AD** check box.

The **AD Domain** row then becomes active and you can enter a value in the **Value** box.

Description	Value
AD Domain (Only used when LDAP Configuration is a Microsoft AD)	

- In the **Value** box of the **The default group** row, specify which OO group is assigned by default to a user that is authenticated by this provider when there is no group matching. If you leave this box empty, only users that have a non-empty OO group set as a result of group mappings are allowed to log on to OO.
- In the **Value** box of the **An internal OO account** row (this is the manager LDAP account), you can choose an internal OO account to be used for directory searches.
- In the **Value** box of the **LDAP search filter that tries to match the user groups** row, 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”

Tom 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 can 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 of the following:

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

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

Make sure you do the following:

- Type each instance of member=CN={1} exactly as it appears above.
 - If you type more than one filter, separate the filters with semicolons (;).
- In the **Value** box of the **List of LDAP contexts containing user groups** row, 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 semicolons (;) to separate the contexts.

For example, suppose that:

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

The following setting makes these groups visible to OO:

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

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

10. In the **Value** box of the **Attribute of any group** row, set the attribute of a group to use as the name of that group when doing group mappings. Usually you can leave the attribute name, but an LDAP schema may use another attribute specific to that schema.
11. In the **Value** box of the **LDAP URL** row, enter the URL or IP address of the top level of the LDAP server, using 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
```

Note: To encrypt communications, you can set OO to communicate with LDAP over SSL by specifying the LDAPS protocol. For information on configuring your system to use the LDAPS protocol, see the *OO Administrator Guide* (`OOAdminGuide.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
```

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

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

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

13. To specify one or more LDAP user context attributes to be used as OO group names, in the **Value** box of the **List of user context attribute names which can be used as groups** row, list the LDAP user context attributes, separating each attribute from the next with a semicolon (;).
14. To specify one or more attributes on which to search for users, provide the attribute names in


the **Value** box of the **LDAP search filter used in the user search** row. For example:

```
&(objectClass=person)(|(sAMAccountName={0})(uid={0})))
```

OR

```
&(objectClass=person)(|(sAMAccountName={0})(userPrincipalName={0})))
```

15. To save your settings, click **Save**.
16. To test the current LDAP authentication settings from within OO, click the **Test** button.
The **Test LDAP Settings** dialog box appears.



17. Type the username and password of an external account that is authenticated by LDAP, and then click **Test**.
You may need to enter the Domain\username or just the username depending on the configuration of the AD Domain.
18. If the test fails, modify the settings and test again.
19. When the LDAP settings test successfully, click the **Save** button in the **LDAP Settings** area.
20. When you see a message to restart the Central service (RSCentral on a Windows system or Central.sh on a Linux system), do so.

To map external LDAP groups to OO groups, see [Mapping external groups to OO groups](#).

Kerberos authentication settings

Kerberos is an authentication-only system, so if you use it you cannot map external groups to OO groups. In order to determine an authenticated user's group memberships, you can use one of the following techniques:

- Like all the other types of authentication providers, you can configure the Kerberos authentication provider to automatically assign a particular OO group (called the default group) to every user it authenticates when no explicit or mapped groups are assigned.
- You can pair the Kerberos provider with an LDAP (standard LDAP or Microsoft AD) provider that can determine authenticated users' group memberships. With this configuration, Kerberos does the authentication, while the LDAP provider supplies the group membership and the mapping to internal OO groups.

Central User Guide

Users, groups, and access control

Note: You must configure the LDAP provider with an LDAP account that has search privileges (the manager account).

- You can assign explicit group membership to Kerberos users. To achieve this, use the explicit user account override mechanism. In versions of OO prior to 9.0, users were tracked only by username; now they are tracked by using the Fully Qualified Name (FQN). In the case of Kerberos, FQN uses the “username@realm” format.

For example, if you used to explicitly assign groups to user “tgage” in the MIRAGE.AD realm by defining a tgage OO (external) account, you will now have to use tgage@MIRAGE.AD to achieve the same result.

The following procedure for authenticating with Kerberos refers to the **Kerberos Settings** area of the **Authentication** tab on the **System Configuration** page.

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

To authenticate with Kerberos

- Click the **Administration** tab, and then click the **System Configuration** tab.
- On the **System Configuration** page, click the **Authentication** tab.
- Scroll down to the **Kerberos Settings** area.

The screenshot shows the 'Administration' tab selected in the top navigation bar. Below it, the 'System Configuration' page is open, with the 'Authentication' sub-tab active. The 'Kerberos Settings' section is expanded, showing a table with configuration options. The 'Kerberos Enabled' checkbox is currently unchecked. The table has two columns: 'Description' and 'Value'. The rows are:

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
LDAP Configuration to lookup the group mappings	

At the bottom right of the settings area, there are 'Test' and 'Save' buttons.

- Select the **Kerberos Enabled** check box.
- In the **Value** box in the **Kerberos 5 configuration file** row, type the name of the file.

The location of the file should be within the OO home directory, and the path should be relative to that directory.

The **Description** column for this setting shows an example of the Kerberos configuration file krb5.conf in the \Central\conf\ subdirectory of the OO home directory.

Caution: If the specified file path for the Kerberos configuration file is invalid, OO

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

6. In the **Value** box in the **KDC host** row, 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 either of the following:

- mydomain.ad:200
- 192.168.5.5:200

7. In the **Value** box in the **Kerberos realm** row, type the domain name of the realm.

For instance, the domain might be MYDOMAIN.AD.

8. In the **Value** box in the **default group** row, specify the default group to which unmapped OO groups or individuals are mapped.

For instance, to assign unmapped groups or individuals only the capability to run flows, you would type **LEVEL_ONE**.


To learn about the default OO groups, see [Managing groups](#).

9. In the **Value** box in the optional **LDAP Configuration to lookup the group mappings** row, you can choose an LDAP provider value from a list the LDAP configurations to use with Kerberos to retrieve the user authorizations and groups.

10. To save your settings, click **Save**.

11. To test the current Kerberos authentication settings from within OO, click the **Test** button.

The **Test Kerberos Settings** dialog box appears.



Test Kerberos Settings

Enter a sample user to Test Authentication with your new settings:
*These credentials are not stored

Username:

Password:

Test Close

12. Type the user name and password of an external account that is authenticated by Kerberos, and then click **Test**.
13. If the test fails, modify the settings and test again.
14. When the authentication settings test successfully, click the **Save** button in the **Kerberos Settings** area.
15. When you see a message to restart the Central service (RSCentral on a Windows system or Central.sh on a Linux system), do so.
16. If your IT organization also authenticates with Active Directory and/or LDAP, complete the procedures in [AD authentication settings](#) or [LDAP authentication settings](#).
17. Finally, because Kerberos authenticates only users, you manually assign external Kerberos users (rather than their groups) to OO groups. For more information, see [Managing users](#).

Managing users

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

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

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

Internal or external users?

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

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

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

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

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

Adding a user

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

To add an individual user

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

The **User Information** dialog appears.



2. Type the user account name.

Note: The colon character (:) is not a valid character in user account names.

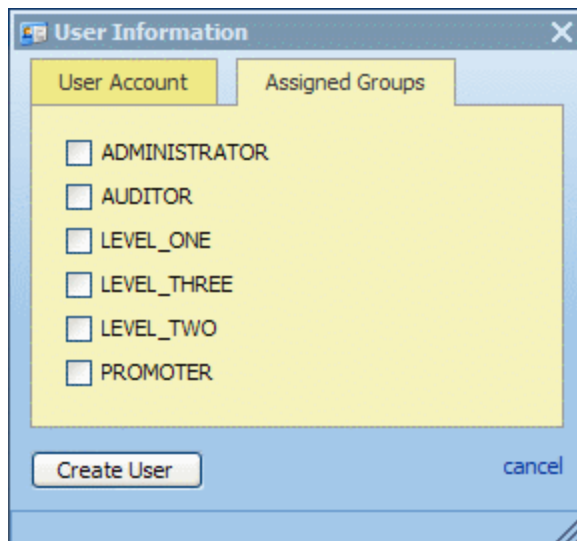
3. If you are creating the account within OO, select the **Internal Account** check box, and then type and verify a password with which the user can log on to OO.

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

The password must be at least six characters long.

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

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



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

Editing a user's account

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

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

To make changes to a user's account

1. On the **Administration** tab's **Manage Users** subtab, click the Edit icon (🔧) next to the name of the user whose account you want to change, then, in the **User Information** dialog, make

changes in the same way that you configured the account when you added the user.

For more information, see [Adding a user](#).

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

Deleting a user

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

To delete a user account

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

Managing groups

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

- Capabilities, or attributes that determine which actions the members of a group can do.

Note: There is not a capability for executing flows. Access to each flow for running it is controlled by its author, who selectively grants the Execute access permission for the flow when he or she creates it.

- Access permissions, which determine which flows, operations (and other parts of flows, such as domain terms) that members of a group can work on.

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

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

Scenario

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

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

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

- The `LEVEL_ONE`, `LEVEL_TWO`, and `LEVEL_THREE` groups are user groups whose rights you define with capabilities and access permissions.

You cannot delete the remaining four default OO groups (`ADMINISTRATOR`, `AUDITOR`, `EVERYBODY`, and `PROMOTER`), which have different specific purposes.

- `ADMINISTRATOR`

The purpose of the `ADMINISTRATOR` group is to have one account that you can use to run and work in Central and Studio in case you temporarily lose the ability to log users in. This group

Central User Guide

Users, groups, and access control

possesses all capabilities and access permissions, so you cannot modify its capabilities or access permissions. You can, however, change the password.

Although you can add members to this group, keep in mind that it is an all-powerful group within OO, so you should assign this account to the fewest people possible.

- **AUDITOR**

As the description indicates, the AUDITOR group might be appropriate for administrators and managers, who should be able to see the data that flows have generated, but who should not necessarily run or author flows. Members of this group have Read permission on all objects and have capabilities that allow them to view flow schedules and create reports.

- **EVERYBODY**


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

- **PROMOTER**

The PROMOTER group has no capabilities, but is able to publish to any Central repository. You cannot change the capabilities of the PROMOTER group.

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

The following schematic shows the promotion of flows.












Promotion of Central repositories

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

Default groups on the Manage Groups subtab

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

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

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

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


Adding groups

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

To add a group

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

The **Group Information** dialog box appears.



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

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

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

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

Adding OO users to groups

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

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

To assign a user to an OO group

1. On the **Administration** tab, click the **Manage Users** subtab.
2. In the row for the user you want to add to a group, click the Edit icon (🔗).

- In the **User Information** dialog box, click the **Assigned Groups** tab, and then specify the groups you want to add the user to.

Mapping external groups to OO groups

Necessary capabilities:

- To enable an external authentication provider, you must be a member of a group that has the `MANAGE_CONF` capability.
- To map an external group to an OO group, you must be a member of a group that has the `MANAGE_GROUPS` capability.

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

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

Notes:

- To map external groups to OO internal groups, you must enable an external authentication provider.

For information on enabling an external authentication provider, see [Using external authentication for Central users](#) and one of the following sections:

- [AD authentication settings](#)
 - [LDAP authentication settings](#)
 - [Kerberos authentication settings](#)
- In this example, you can define what the `LEVEL_ONE`, `LEVEL_TWO`, and `LEVEL_THREE` groups can do by assigning capabilities to the group. So to enable the LDAP group Quality Assurance to test flows, you might assign the `LEVEL_TWO` group in OO the `SCHEDULE`, `MANAGE_RUNS`, `VIEW_SCHEDULES`, and `HEADLESS_FLOWS` capabilities.

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


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

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

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

To map external groups to OO groups

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



3. In the text box, type the name(s) of the external group or groups whose members you want to be members of this OO group.

For instance, to map the AD or LDAP "Network Specialist" group to the OO LEVEL_THREE group, type **Network Specialist** in the text box.

4. Click **Update Group**.

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

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

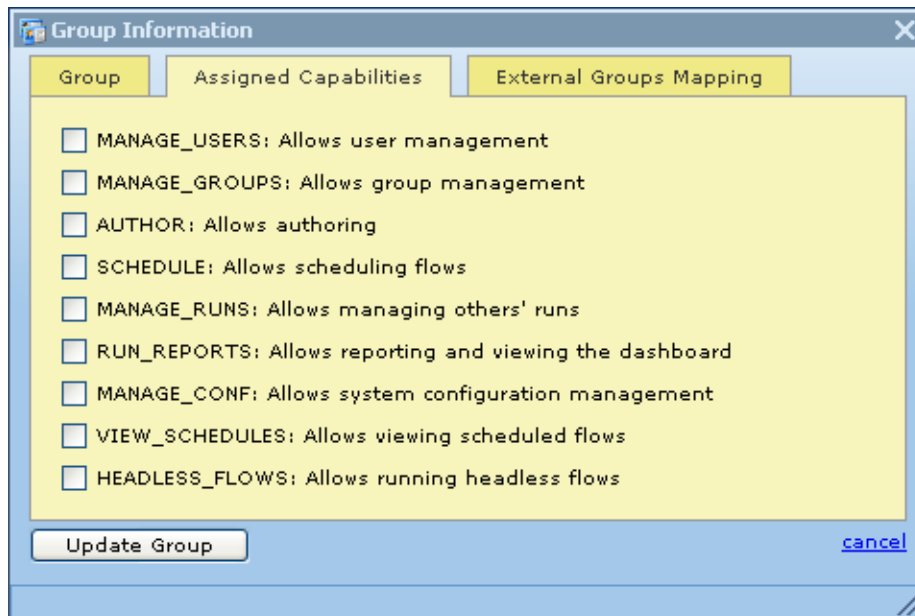
Changing a group's assigned capabilities and description

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

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

To change the capabilities assigned to a group

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



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

Note: Notes:

- You cannot change the capabilities for the `ADMINISTRATOR`, `AUDITOR`, `EVERYBODY`, or `PROMOTER` group. For information on the intended uses of these groups, see the Operations Orchestration *Administrator's Guide* (OOAdminGuide.pdf).
- By default, the groups `LEVEL_ONE`, `LEVEL_TWO`, and `LEVEL_THREE` have no capabilities, so you must assign them some.
- For information on capabilities and the difference between them and access permissions to objects, see the *OO Administrator Guide* (OOAdminGuide.pdf).

4. If your conception of the group changes after you change the capabilities you grant it, you may want to click the **Group** tab and change the group's name and description to make them more descriptive.
5. When you've finished working here, click **Update Group**.

In addition, the following tasks must be executed outside the Central application. For more information on these tasks, see the *OO Administrator Guide* (OOAdminGuide.pdf).

Central User Guide

Users, groups, and access control

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

Deleting groups










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

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

To delete a group

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

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

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

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

A run is a single execution of a flow. At any given time, there may be many flow runs running in Central. While flow runs are still in progress, you can do the following:

- View a list of all active flow runs and view information about them.
- Narrow the list of flow runs shown by [filtering them](#).
- [View flow run details](#).
- Interrupt flow runs.
- [Resume flow runs](#).
- [Delete flow runs](#).
- [Reassign flow run ownership](#).
- [Watch the progress of selected flow runs](#).
- [View flow run properties](#).

All of these tasks, except interrupting flow runs, are performed on the **Current Runs** tab.

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



Select All	Name	Run ID	Type	Step	Started	Step Name	Step Start Time	Previous Step Name	State	User
<input type="checkbox"/>	Connectivity_Test_6	6	Interactive	0	Dec 30, 2011 3:42:31 PM	is SQL Connected	Dec 30, 2011 3:46:1...	Connectivity Test	Stopped	admin
<input type="checkbox"/>	URL_Test_5	5	Interactive	9	Dec 30, 2011 12:26:04 PM	Resolved : success	Dec 30, 2011 1:24:3...	Display check	Input Required	admin
<input type="checkbox"/>	Windows_Health_Check_2	2	Interactive	3	Dec 30, 2011 12:14:07 PM	Resolved : success	Dec 30, 2011 12:22:...	Evaluate Expression	Input Required	admin
<input type="checkbox"/>	Windows_Health_Check_1	1	Interactive	3	Dec 30, 2011 11:59:39 AM	Resolved : success	Dec 30, 2011 12:14:...	Evaluate Expression	Input Required	admin

By default, the flow run list in the **Current Runs** tab is sorted by run ID, starting with the most recent run. You can reverse this order by clicking the down arrow in the **Run ID** title row. You can also sort the flow run list by any column header by clicking that header.

The times given on the **Current Runs** tab are given in the time zone for the Central client, whose relationship to Greenwich Mean Time (GMT) is shown in the Central web application's title bar.

Refreshing the Current Runs tab

Because the flow run environment is constantly changing, the **Current Runs** tab needs to be frequently refreshed. You can use the **Auto-Refresh** control to refresh the tab and update the flow run information automatically every 15 or 30 seconds.

On the other hand, you may wish to refresh the tab manually.

To set the Auto-Refresh control

1. Locate the **Auto-Refresh** control area, which looks like the following.

Auto Refresh: OFF | 15 | 30 seconds 

2. To refresh the **Current Runs** tab manually only:
 - a. Select **OFF**.
 - b. To refresh the tab, in the **Filter Panel**, click **Apply**.
3. To refresh the **Current Runs** tab automatically, in the **Auto-Refresh** control area, click either **15** or **30 seconds**, according to how frequently you want it to refresh.

Note: The red indicator light in the **Auto-Refresh** control flashes when it communicates with the Central server in the course of refreshing.

Filtering the flow runs to view

If you have the `MANAGE_RUNS` capability, the **Current Runs** tab lists all of the flow runs currently running in Central (if you do not have the `MANAGE_RUNS` capability, you will only see your active flow runs). This is probably more information than you want to view at one time.

To narrow the list down to a more manageable size, you can:

- Use the **Filter Panel** on the **Current Runs** tab to create a filter that selects only those flow runs you want to view at the moment.

For more information, see [Filtering flow runs using the Filter Panel](#).

- Apply additional filtering with column headers.

For more information, see [Filtering flow runs with column headers](#).


Note: Filtering only affects the flow runs in the **Current Runs** tab, not those in the **Watched Runs** pane.

Filtering flow runs using the Filter Panel

To filter flow runs using the Filter Panel

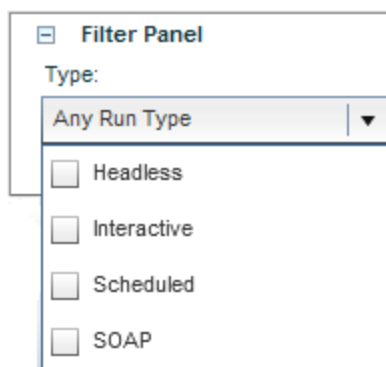
1. Click the **Current Runs** tab.

The **Filter Panel** on the **Current Runs** tab has four drop-down list boxes you can use to specify criteria for filtering your view of the currently running flows.



You can filter flow runs using any combination of the **Type**, **State**, **Run Identifier**, and **User** criteria, which are combined with an AND operator to produce results. So, entering a value of **bsmith** for the **User** criterion and selecting a value of **headless** for the **Type** criterion, the resulting list shows only headless runs whose user is bsmith.


The **Type** and **State** drop-down lists allow you to choose one or more values from a drop-down list.



If you select multiple values in a filter's drop-down list, they are combined with an OR operator. For example, selecting the **headless** and **interactive** values in the **Type** drop-down list shows runs that are either headless (that is, that are started from outside Central or by the Scheduler) or interactive (that is, are started by a user in OO Central).

2. To add **Type** criteria to your filter, in the **Type** drop-down menu, do one of the following:
 - To select flow runs of any type listed below, leave the default of **Any Run Type** selected.
 - Select one or more of the following:
 - **Headless** - Selects flow runs that are started outside of Central using the RSFlowInvoke.exe or JRSFlowInvoke.jar utility.
 - **Interactive** - Selects flow runs that are started from the Central Flow Library tab.
 - **Scheduled** - Selects flow runs that are started by the Scheduler.
 - **SOAP** - Selects flow runs that are started by the SOAP Web service.
3. To add **State** criteria to your filter, in the **State** drop-down menu, do one of the following:
 - To select flow runs of any state listed below, leave the default of **Any Run State**.
 - Select one or more of the following:
 - **NEW** – Selects flow runs that have not yet completed their first step.
 - **RUNNING** – Selects flow runs that are currently running.

- **STOPPED** – Selects flow runs that are stopped.
 - **STOP PENDING** – Selects flow runs that are pending being stopped.
 - **CANCELING** – Selects flow runs that are being canceled.
 - **HANDED OFF** – Selects flow runs that have been handed off.
 - **HAND OFF PENDING** - Selects flow runs that are pending being handed off.
 - **INPUT REQUIRED** – Selects flow runs that require user input to continue.
 - **CREDENTIALS REQUIRED** – Selects flow runs that require the user to enter credentials to continue.
 - **ORPHANED** – Selects flow runs that are running when Central fails.
4. To add **Run identifier** criteria to your filter, in the **Run Identifier** drop-down menu, do one of the following:
- To select flow runs with all run names or IDs, leave the default of **Any Run Identifier**
 - Select **With Run Identifier...**
- If you select **With Run Identifier...**, the **Runs with the run ID(s), run flow name(s) or flow UUID(s)...** box opens.



Runs with the run ID(s), run flow name(s) or flow UUID(s)...

Enter run ID(s), run flow Name(s) or flow UUID(s) on their own line, or separated by commas.

OK Cancel

- Do the following:
 - i. Enter the run IDs (not history IDs), flow run names (usually the flow names), or flow UUIDs of the flow runs you want to view, each on its own line or on one line separated by commas.
- Notes:**
- Run IDs are used to filter active runs while history IDs are used by the reports (after a run has finished executing).
 - If you filter on a flow names, you can enter a partial flow run name; this returns runs with flow names that start with the partial name provided.
- ii. Click **OK**.

5. To add **User** criteria to your filter, in the **User** drop-down menu, do one of the following:

- To select only those runs started by you, leave the default of **My Runs**.
- Select one of the following:
 - **Any User** – Selects flow runs started by any user
 - **Runs By...** – Selects flow runs started by specified users

If you select **Runs By...**, the **Run executed by users...** box opens.

- Do the following:
 - i. Enter the usernames of those users whose flow runs you want to view, each on its own line or on one line separated by commas.

Note: Usernames may include domains—“domain\username” for Windows systems or “username@domain” for Linux systems. For example, domain1\bsmith or bsmith@domain1.
 - ii. Click **OK**.

Note: Filtering based on usernames matches only exact names, not partial names.

6. Click **Apply**.

The selected flow runs appear in the **Current Runs** pane.

Filtering flow runs with column headers

You can further filter the list of flow runs obtained with the Filter Panel by using the **Name**, **Run ID**, **Type**, **State**, and **User** column header filters shown below.

Select All	Name	Run ID	Type	Step	Started	Last Modified	State	User
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>				<input type="text"/>	<input type="text"/>

Note: The **Started** and **Last Modified** times are given in the time zone for the Central client, whose relationship to Greenwich Mean Time (GMT) is shown in the title bar of the web page.

To filter flow runs using the column header filters

- In a column header filter, type values that further filter the list of flow runs already obtained.

Column header filters can filter based on partial inputs. For example, typing **service** in the **Name** column filter header matches flow run names such as "Start Service" and "Service Status".

You can also refine your search by using standard regular expressions.

For example, typing **^service** in the **Name** column header filter matches only those flow run names that begin with the word "service", such as "Service Status".

Select All	Name	Run ID	Type	Step	Started	Last Modified	State	User
<input type="checkbox"/>	Get							
<input type="checkbox"/>	Get Database ... ?	7	Interactive	0	03/08/10 12:53	03/08/10 12:53	Stopped	admin
<input type="checkbox"/>	Get Process Info ?	6	Interactive	0	03/08/10 12:52	03/08/10 12:53	Stopped	admin

Resetting filters to their defaults

To reset all of the filters to their default settings

- Click the **Reset** button in the **Filter Panel**.

Filter Panel

Type: State: Run Identifier: User:

Manage Filters

Saving filters

You can save the current filter values in the **Filter Panel**, and then load them later from a list of saved filters.

To save a filter

1. On the **Current Runs** tab, click the **Manage Filters** link.


Filter Panel

Type: State: Run Identifier: User:

[Manage Filters](#)

Manage Filters link


The **Manage Filters** box opens.



Manage Filters box

2. Click **Save As**.

The **Manage Filters** box expands to display the **Filter name** input.



Manage Filters box expanded

3. Enter a name for the filter and then click **Save**.

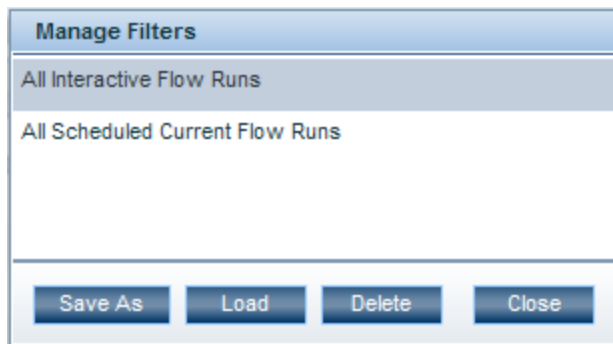
Note: The **Filter Panel** will display the saved filter name.

Loading a saved filter

To load a saved filter

1. On the **Current Runs** tab, click the **manage filters** link.

The **Manage Filters** box opens, displaying any saved filters.



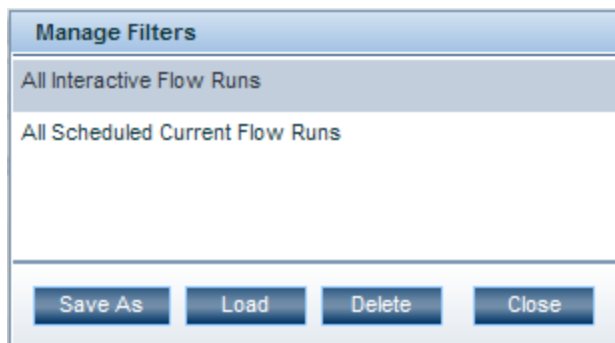
2. Double-click a filter name or select a filter name and then click **Load**.

Deleting a saved filter

To delete a saved filter

1. On the **Current Runs** tab, click the **manage filters** link.

The **Manage Filters** box opens, displaying any saved filters.



2. Select a filter name, and then click **Delete**.

Viewing the details of a flow run

If a flow run has not yet completed, you can view a flow run's details on the **Current Runs** tab. If it has completed, you can see its details either on the **Reports** tab or in the **Watched Runs** panel of the **Current Runs** tab.

The **Current Runs** tab initially shows, for each run that is started but has not completed, information that tells you which step is currently being executed in the run, what the last step was, etc.

Select All	Name	Run ID	Type	Step	Started	Step Name	Step Start Time	Previous Step Name	State	User
<input type="checkbox"/>	Windows Health Check_308	308	Interactive	32	Jan 02, 2012 8:45:33 PM	Display Message: ...	Jan 02, 2012 9:02:4...	Display?	Stopped	admin
<input type="checkbox"/>	Windows Health Check_136	136	Interactive	0	Jan 02, 2012 12:14:21 PM	Get Host Name	Jan 02, 2012 12:52:...	Get System Uptime	Stopped	admin
<input type="checkbox"/>	Windows Health Check_104	104	Interactive	5	Jan 02, 2012 3:26:03 AM	Get Host Name	Jan 02, 2012 4:47:3...	Get System Uptime	Stopped	admin

The **Reports** tab shows detailed information about flow runs, particularly what has taken place in each step. You can access the detailed report by drilling down in the **Reports** tabs.

To access a current flow run's details from the Current Runs tab

- On the **Current Runs** tab, click the run name in the **Current Runs** pane or the **Watched Runs** pane whose details you want to view.

The **Reports** tab shows information for the steps that have completed so far, with the response for each step and the message that describes the result for that step.

Results > One Run Instance

Run Id: 104
Run Name: Windows Health Check_104

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

Basic | **Advanced**

Step #	Step Name	Response	Message	Call Path
1	Ping Target System	✓	16.59.62.4 successfully pinged.	Windows Health Check > Ping Target System
2	Get System Uptime	✓	16.59.62.4 has been running for SystemUpTime 1423319 seconds.	Windows Health Check > Get System Uptime

To access the details of completed flow runs from the Reports tab

1. On the **Reports** tab, click the name of the flow of interest.

The **Reports** tab changes to show runs that have been completed for the flow.

Run Id	Run Name	Start Time	Executed by	Time To Run	Result Type	Number of Steps	Last Modified by	Revision	ROI Value
461	Windows Health Check_461	Jan 2, 2012 9:02:46 PM	admin	123.233	Resolved	36	admin	186	0
140	Windows Health Check_140	Jan 2, 2012 5:04:35 PM	admin	42.046	Resolved	39	admin	186	0

2. Click the run ID of the run of interest.

Among other information, the **Reports** tab displays the responses obtained and information generated by the run's steps.

Step #	Step Name	Response	Message	Call Path
1	Ping Target System	✓	16.59.62.4 successfully pinged.	Windows Health Check > Ping Target System
2	Get System Uptime	✓	16.59.62.4 has been running for SystemUpTime 1482746 seconds.	Windows Health Check > Get System Uptime
6	Get Host Name	✓	16.59.62.4 believes its name to be Name Microsoft Windows Server 2003 Enterpris...	Windows Health Check > Get Host Name
10	Get Processor Count/Status	✓	16.59.62.4 has 2 processors. They were all online. Their loads were: 0,49% . They had 5...	Windows Health Check > Get Processor Count/Status
11	Free Physical Memory	✓	16.59.62.4 has FreePhysicalMemory 2433436 kilobytes of free memory .	Windows Health Check > Free Physical Memory
15	Get Running Process List	✓	16.59.62.4 has the following processes running: Name [System Idle Process;System;...	Windows Health Check > Get Running Process List
19	Get FileSystem Information	✓	File system information retrieved: FreeSpace,Name,Size " ",A;" " 69442519040,C,8...	Windows Health Check > Get FileSystem Information
23	Check for Malfunctioning Devices	✓	The following devices were malfunctioning, if any:	Windows Health Check > Check for Malfunctioning Devices
27	Get Network Interface Errors	✓	The network interface had no errors	Windows Health Check > Get Network Interface Errors
31	Create System Health Report	✓	Health report created: Server Health Report for server 16.59.62.4: Summary Of ...	Windows Health Check > Create System Health Report
32	Notify	✓	Health report sent.	Windows Health Check > Notify

What happened in a given step is described in the **Message** column. The information provided there is determined by what the author of the flow. When the flow is created, the author specifies in each flow step's message what information will be provided to the user in that step's message.

3. To see more of the information generated by the flow run, click the **Advanced** tab. The information presented on the **Advanced** tab includes the following for each step in the flow run:

- Step name
- Step number
- User ID
This is the user account name
- Parent flow
This is the flow from which the step was created.
- Start time and end time of the step


These times are given in the time zone of the client machine from which the flow was run.

- Execution time
This is how long the step took to run, in seconds.
- Description
- Response
- Bound inputs and their values (unless encrypted)
Bound inputs are the inputs and their assigned values.
- Recorded bound inputs
- Result
- Description
- The ROI value of the transition, if any has been defined for it.

Taking control of a flow run

You can take control of flow runs, to resume, delete, or otherwise work with them. Note that you can only take control of one run at a time.

To take control of a flow run

1. Click the **Current Runs** tab.
2. In the **Current Runs** or **Watched Runs** pane, select the checkbox in the leftmost column of the row for the flow run you want to take control of.
3. Click the  button.

The flow run is placed into a Stopped state, and the flow's preview page appears, with the flow run state showing as Stopped.

4. According to how you want to resume the flow run, click either of the following to continue the run in one of the run modes:
 - **Next Step**, to continue the run in the Guided Run mode
 - **Run All**, to complete the run without further intervention.

For more information about run modes and running a flow in each of those modes, see respectively [Types of flow runs](#) and [Running a flow](#).

Notes:

- To resume a run that you did not start and that has not been re-assigned or handed off to you, you must be a member of the ADMINISTRATOR group or of a group that has the MANAGE_RUNS capability.
- The only runs that you should resume are those that have been interrupted or handed off to you. Clicking **Resume** for a run that is currently running transfers ownership of the run from

the user who started it to you. This causes an error for the user whose active run you resumed. In addition, other problems may result. Active runs include those whose state is NEW, RUNNING, INPUT REQUIRED, or CREDENTIALS REQUIRED. For information on reassigning ownership of a run, see [Reassigning ownership of a run](#).


- If you resume a run that is already running (this includes headless runs), it will go into a STOPPED state, and you must click either **Next Step** or **Run All** in the flow graph.

Resuming a flow run

Notes:

- **Required capabilities:** To resume a run that you did not start and that has not been re-assigned or handed off to you, you must be a member of the ADMINISTRATOR group or of a group that has the MANAGE_RUNS capability.
- The only runs that you should resume are those that have been interrupted or handed off to you. Clicking **Resume** for an active run (one that is currently running) transfers ownership of the run from the user who started it to you. This causes an error for the user whose active run you resumed. In addition, other problems may result. Active runs include those whose state is NEW, RUNNING, INPUT REQUIRED, or CREDENTIALS REQUIRED. For information on reassigning ownership of a run, see [Reassigning ownership of a run](#).

To resume a flow run

1. Click the **Current Runs** tab.
2. In the **Current Runs** or **Watched Runs** pane, select the checkbox in the leftmost column of the row for the flow run you want to resume.
3. Click the  button.

The flow run is placed into a Stopped state, and the flow's preview page appears, with the flow run state showing as Stopped.

4. According to how you want to resume the flow run, click either of the following to continue the run in one of the run modes:
 - **Next Step**, to continue the run in the Guided Run mode
 - **Run All**, to complete the run without further intervention.


For more information about run modes and running a flow in each of those modes, see respectively [Types of flow runs](#) and [Running a flow](#).

Deleting flow runs

You can delete one or more flow runs before they complete execution. When you delete a flow run, the details of that flow run are no longer available to you, including the flow run history accumulated up to that point.

To delete flow runs

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



2. In the leftmost column, select the check boxes of the flow runs you want to delete.

OR


To select all the current runs, select the check box under **Select All** in the heading row.



The check boxes for all of the flow runs are selected.

3. Click the **Delete** icon ().

The **Delete the following item(s)...** box opens.



4. Click **OK**.

Reassigning ownership of flow runs

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

Notes:

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

- Reassigning a run to a different user causes the run to stop if it was not already in a STOPPED state.
- The original user receives a notification about the change of ownership.
- Reassigning a run to the same user should have no effect.

To reassign ownership of flow runs


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

Select All	Name	Run ID	Type	Step	Started	Step Name	Step St
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>				
<input type="checkbox"/>	Iterative CPU Diagnostic_9	9	Interactive	1	Jan 04, 2012 1:27:09 AM	Iterator	Jan 04, 20
<input type="checkbox"/>	Restart Service_8	8	Interactive	0	Jan 04, 2012 1:26:17 AM		Jan 04, 20
<input type="checkbox"/>	vWindows Health Check_7	7	Interactive	5	Jan 04, 2012 1:23:41 AM	Get Host Name	Jan 04, 20

2. In the leftmost column, select the check boxes of the flow runs you want to reassign.

OR

To select all the current runs, select the check box under **Select All** in the header row of the table.

3. Click the **Reassign** button ().

The **Reassign the following item(s)...** box opens.

Reassign the following 1 item(s)...

Name	Run ID	History ID	User
Restart Service_8	8	8	admin

Assign run(s) to:

4. In the **Assign run(s) to** box, enter the username of the person to which you want to reassign the flow run or runs.
5. Click **OK**.

Watching runs

You can select a group of flow runs to display in the **Watched Runs** pane so that you can closely observe their progress. The **Watched Runs** pane displays the same information for flow runs as the **Current Runs** pane.

Notes:

- The selected runs to be watched continue to appear in the **Current Runs** pane.
- The watched runs continue to appear in the **Watched Runs** pane after they finish running.
- The **Watched Runs** pane can display a maximum of 100 flow runs at any time.
- By default, the flow run list in the **Watched Runs** pane is sorted by run ID, starting with the most recent run. You can reverse this order by clicking the down arrow in the **Run ID** column header. You can also sort the flow run list by any column header by clicking the header.

To watch runs

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

The runs that have started but not finished are listed in the **Current Runs** pane.

2. In the left-most column, click the box of each run that you want to watch.

OR

To select them all, click the box at the top of the column, under **Select All**.

Select All	Name	Run ID	Type	Step	Started	Step Name	Step Start Time
<input type="checkbox"/>							
<input checked="" type="checkbox"/>	Connectivity Test_266	266	Interactive	1	Dec 31, 2011 10:58:30 PM	Is lossThreshold va...	Dec 31, 2011 11:01:...
<input type="checkbox"/>	vWindows Health Check_254	254	Interactive	5	Dec 31, 2011 10:53:49 PM	Get Host Name	Dec 31, 2011 10:58:...
<input checked="" type="checkbox"/>	How do I: Format with HTML_209	209	Interactive	1	Dec 31, 2011 10:38:07 PM	Write to File	Dec 31, 2011 10:38:...
<input type="checkbox"/>	vWindows Health Check_178	178	Interactive	1	Dec 31, 2011 9:41:34 PM	Get System Uptime	Dec 31, 2011 10:05:...

3. Click the **Watch** button.

The selected runs appear in the **Watched Runs** pane, which has the same layout as the **Watched Runs** pane.

Select All	Name	Run ID	Type	Step	Started	Step Name	Step Start Time	Previous Step Name	State
<input type="checkbox"/>									
<input type="checkbox"/>	Connectivity Test_266	266	Interactive	1	Dec 31, 2011 10:58:30 PM	Is lossThreshold va...	Dec 31, 2011 11:37:...	Is latencyThreshold vall...	Stopped
<input type="checkbox"/>	How do I: Format with HTML_209	209	Interactive	1	Dec 31, 2011 10:38:07 PM	Write to Is lossThreshold valid? 2011 10:38:...		HTML String Formatter	Orphaned

4. Delete or reassign any flow run in the **Watched Runs** pane the same way you would for a flow run in the **Current Runs** pane. See [Deleting a flow run](#) and [Reassigning ownership of a flow run](#) for instructions.
5. To see the report of the steps that have been executed so far, click the run name.

To remove flow runs from the Watched Runs pane

1. In the **Watched Runs** pane, select the check boxes next to the runs you want to remove.

OR

To remove all of the flow runs in the **Watched Runs** pane, select the check box in the **Select All** column header.

2. Click the **Unwatch** button.

Viewing flow run properties


From the **Current Runs** tab, in the **Current Runs** and **Watched Runs** panes, you can see the following properties for a flow run:

- The name of the flow
- The globally unique identifier (UUID) of the flow run
- Specifics about the flow that are available in the flow description in the Flow Library or, for authors, on the flow's **Description** tab in OO Studio.

To view a flow run's properties

1. On **Current Runs** tab, in either of the **Current Runs** and **Watched Runs** panes, in the row of the flow run of interest, select the check box in the left-most column.

You can only view the properties for a single flow run at a time.

2. In the toolbar for the pane you're working in, click the **Properties** button ().

A message box appears, displaying the information described above.

Managing Central nodes

An OO Central node is a single installation of OO Central (whether it's in a clustered or non-clustered OO deployment). There may be instances when you may need to manage an individual Central nodes' state. Using the Node Administration page, you can take nodes on- and offline, view the nodes' usage and status, and connect directly to them.

Required permissions: You must have `MANAGE_CONF` permissions to view the **Node Administration** page.

To view the OO Node Administration page

1. In Central, click the **Administration** tab.
2. On the **Central Administration** page, click the **Node Administration** tab.

This opens the **Node Administration** page.

Auto Refresh: OFF | 15 | 30 seconds

Central Node Management

Action	Node	Status	Active Runs	Studio Users	Central Users
Take Offline	TC-QA4ND2.battleground.ad:8443	ONLINE	0	0	0
Take Offline	TC-QA4ND1.battleground.ad:8443	ONLINE	0	0	1

Counters cached for 15 seconds.

Remote Action Service (RAS) Infrastructure View

Name	URL	Connection Status	Version	Platform
RAS_Operator_Path	https://TC-QA4ND1.battleground.ad:9004/RAS/services/RCAgentService	Connecting from all Central Nodes	7.80.0.0	Windows 2003

Connection Status cached for 90 seconds.

The information in these panes is static, but you can refresh a pane by clicking the **Refresh** link in the pane's toolbar, or by clicking **15** or **30** in the **Auto Refresh** link at the top of the page. Click **15** for auto refreshing every 15 seconds or **30** for every 30 seconds.

The **Node Administration** page has two panes:

- The **Central Node Management** pane, where you can:
 - Take nodes offline
 - Bring nodes online
 - View the number of active runs, Central users, and Studio users on each node

- The **Remote Action Service (RAS) Infrastructure** pane, where you can view information about your RAS infrastructure, including the connection status, version, and platform of each RAS defined in the repository. See [Viewing RAS connections](#) for more information.

Viewing the Central nodes


The **Central Node Management** pane of the **Node Administration** page displays all of the Central nodes in your cluster. For each node, the pane also displays counters for the number of active runs, Studio users, and Central users on it. This provides you with a snapshot of each node's usage.

Note: When you take a node offline, the counters should decrease as users with existing Central and Studio sessions stop using the node or when their sessions expire.

The **Central Node Management** pane of the **Node Administration** page provides a way to see more detailed information about nodes.

To view detailed information about nodes

1. In the toolbar of the **Central Node Management** pane, click the **Show Node Details** link.
This expands the area to display the **Link to node** link and the node ID.
2. You can log in directly to a node by clicking the **Link to node** link for the node.



Action	Node	Status	Active Runs	Studio Users	Central Users
Take Offline	USER1.amecorp.net:8443 Link to node... Node ID: central-USER 1.amecorp.net-9e5103cb-f210-4c24-8b7e-8917750fc9c	ONLINE	0	1	1

The node opens in a new window.

Taking nodes offline

You can take nodes offline on the **Central Node Management** pane of the **Node Administration** page.

Taking a Central node offline affects the OO components as follows:

- The Scheduler
The Scheduler is taken offline, so that the Scheduler does not attempt to schedule new runs for the offline Central node.
- Flow runs
New flow runs are blocked.
- Headless and SOAP requests
New headless and SOAP requests are blocked. If the node receives a new headless or SOAP request, a message appears, stating that the node is offline or showing an HTTP 503 error code.

Existing SOAP sessions continue to work until the session expires; however, new flow runs are still blocked.

- Studio

New Studio sessions are blocked. Users who try to log on to Studio on an offline node receive a message stating that the node is offline.

Studio users who have an current session open when the Central node goes offline receive a message telling them that the repository is going offline. The message asks them to save their work and log out.

Studio users can continue to work in this repository until their session expires or until the Central service or the node is shut down.

- Central sessions

New Central sessions are blocked. Users who try to log on to an instance of Central that is on an offline node receive a message stating that the node is offline.

Central users who have an existing session can continue to work until they log out, the session times out, or the Central Service or the node is shut down.

- Load Balancer

If you are using the HP OO Load Balancer and load balancer management is enabled, the specific node on the load balancer is put in stand-by mode and the Load Balancer stops routing new traffic to that Central node. Existing sessions initiated by the load balancer remain active until you log out, the session times out, or the Central service or the node is shut down.

If you are using a third-party load balancer, the load balancer's behavior depends on its features and how the load balancer management flow is configured to administer it. For more information on load balancer management flows, see [Integrating a load balancer with OO node management](#).

To take a Central node offline

- Click the **Take Offline** button in the **Action** column of the node you want to take offline.

The node's OFFLINE status appears in the **Status** column, and the **Action** column displays the **Bring Online** button.

Bringing nodes online

You can bring a node online either on the Central Node Management pane of the Node **Administration** page or by starting Central.

When you place a Central node online, it does the following:

- Allows new requests to come through (for instance, headless and SOAP requests, Studio sessions, and Central sessions).
- Brings the attached Scheduler online so that flows can be scheduled on the Central node.
- If you are using the HP OO Load Balancer and load balancer management is enabled, the node on the load balancer is resumed and starts routing new traffic to the Central node. If you are using a third-party load balancer, the load balancer's behavior depends on its features and how

the load balancer management flow is configured to administer it. See [Integrating a load balancer with OO node management](#) for more information on load balancer management flows.

To bring a Central node online

- Click the **Bring Online** button in the **Action** column of the node you want to bring online.

OR

Restart Central.

Integrating a load balancer with node management

OO deployments can be set up with load balancing. Read the *High Availability Guide (Clustering_LoadBalancingGuide.pdf)* for details. In most cases where load balancers are used, when a Central node's state changes, it is important that these state changes be reflected in the load balancer so that it routes traffic correctly.

To synchronize the state of the load balancer with OO node state changes, you need a load balancer management flow to manage the load balancer's state. How you create the load-balancer state management flow depends on the type of load balancer that you use:

- The HP OO Load Balancer, in which case you can use the HP Load Balancer admin flow
- A third-party load balancer, for which you must create your own load-balancer management flow

The load-balancer management flow must synchronize the state of the load balancer with OO state changes. Such a flow requires that the load balancer you are using must provide an integration point for OO (for instance, Web services or Java/.Net API) or some method to communicate with its management console (for instance, Client, SSH, or HTTP). Once there is a way to communicate with the load balancer, you can use existing OO content and OO wizards, or build your own OO operations for the load balancer management flow.

To learn how to create OO operations and flows, see the *OO Studio Authoring Guide (Studio_AuthorsGuide.pdf)*.

When the flow is available for use in Central, then, on the **Clustering Settings** subtab on the **System Configuration** tab:



1. Enable load balancer management.
2. Specify the load balancer management flow in Central's system configuration.

For more information on these two tasks, see [Enabling load balancer management for a Central cluster](#).

Note: If you are using the HP Load Balancer and the provided HP Load Balancer admin flow, when you set up the Load Balancer you must use the same URLs for the nodes that appear in the **Node** column of the **Central Node Management** pane in the **Node Administration** page.

Central User Guide

Managing Central nodes

Central Node Management					
Refresh Show Node Details					
Action	Node	Status	Active Runs	Studio Users	Central Users
Take Offline	TC-QA4ND2.battleground.ad:8443 ↗	 ONLINE	0	0	0
Take Offline	TC-QA4ND1.battleground.ad:8443 ↗	 ONLINE	0	0	1

Counters cached for 15 seconds.

These URLs are passed from OO to the load balancer management flow. If the URLs do not match the nodes in the load balancer, the flow fails and the load balancer continues with its state unchanged.

Viewing RAS connections

You can view all of your RAS connections and details about those connections on the **Node Administration** page.

To view your RAS connections


1. In Central, click the **Administration** tab.
2. On the **Administration** page, click the **Node Administration** tab.

Every RAS defined in the repository appears in the **Remote Action Service (RAS) Infrastructure View** pane along with information such as its connection status and RAS-specific information like its version and platform. If any Central node is unable to reach a RAS, this view indicates that there are connection issues and allows the user to see details of the error causing the connection issue.

3. To view the connection details of a specific RAS connection, click the connection status in the **Connection Status** column.

OR

To see the details of all of the RAS connections, in the toolbar of the **Remote Action Service (RAS) Infrastructure View** pane, click the **Show Connection Details** link.



The screenshot shows the 'Remote Action Service (RAS) Infrastructure View' pane. It includes a toolbar with 'Refresh' and 'Hide Connection Details' buttons. Below is a table with columns: Name, URL, Connection Status, Version, and Platform. The table contains one entry for 'RAS_Operator_Path' with a URL pointing to a RAS service. The 'Connection Status' column shows a green checkmark and the text 'Connecting from all Central Nodes' with a dropdown arrow. Below this, two connection attempts are listed, both showing 'Connection success' with their respective URLs and update times. The 'Version' is 7.80.0.0 and the 'Platform' is Windows 2003. A note at the bottom right states 'Connection Status cached for 90 seconds.'.

Name	URL	Connection Status	Version	Platform
RAS_Operator_Path	https://TC-QA4ND1.battleground.ad:9004/RAS/services/RCAgentService	✓ Connecting from all Central Nodes https://TC-QA4ND1.battleground.ad:8443 updated: 2009-11-12 14:39:29.761 Connection success https://TC-QA4ND2.battleground.ad:8443 updated: 2009-11-12 14:39:48.761 Connection success	7.80.0.0	Windows 2003

The RAS connectivity information is gathered periodically based on a timer, so RAS availability changes may not appear instantaneously.

Managing your runtime environment with overrides

OO customers may have multiple OO deployments serving various purposes such as development or production, or separated geographically. In such cases, it is common to share repository content between the deployments. Sharing these repository items can become challenging and unwieldy if some of them contain data that is specific to a particular OO deployment. Configuration items, step inputs, and flow inputs are examples of items in operations and flows that may be specific to an OO deployment.

The OO Runtime Environment feature allows you to override certain OO configuration item types previously created in Studio—Remote Action Services (RASes), system accounts, and system properties—in individual OO deployments (see the OO Studio Authoring Guide (Studio_AuthorsGuide.pdf) to learn how to create RASes, system accounts, and system properties). Doing so keeps the overridden values outside of the OO repository, substituting them at runtime. Separating the deployment-specific data simplifies the publish/update process.

To avoid having deployment-specific data in OO repositories, users can design their operations and flows to use system properties, system accounts, and RAS operators where appropriate. Then, when you publish to another OO deployment, you publish the operations, flows, and related items, and then update the runtime environment overrides to change deployment-specific items.

For example, let's say that you have two OO deployments (Deployment 1 and Deployment 2) and each deployment uses its own LDAP server (ldaphost1.domain.com and ldaphost2.domain.com). You need a flow in both deployments that returns a list of user attributes from its own LDAP.

Using Studio, you create a flow in Deployment 1 named Get LDAP User Attributes that uses a system property named ldap_host which points to Deployment 1's LDAP server ldaphost1.domain.com. Since you want to use the same functionality in all of your OO deployments, you publish the Get LDAP User Attributes flow and related content (the ldap_host system property) from Deployment 1 to Deployment 2. The problem is that the ldap_host system property now on Deployment 2 is directing the flow to ldaphost1.domain.com when you need it to be directed at ldaphost2.domain.com.

This can be solved by one of the following methods:

- Manually editing the Get LDAP User Attributes flow in Deployment 2.
- By adding a runtime environment override for ldap_host on Deployment 2 that sets the value to ldaphost2.domain.com. The figure below illustrates this solution and includes a use case for RAS overrides.


In the RAS use case shown below, if you do not manually update the RAS object in the Deployment 2 repository or add a RAS environment override, Central tries to connect to the Deployment 1 RAS (ras1.domain.com) when the flow is executed.

When you subsequently edit the flow, you will simply need to publish the changes from Deployment 1 to Deployment 2 (unless you are adding new deployment-specific information that needs to be

Central User Guide

Managing your runtime environment with overrides

edited or overridden). Using the environment override feature greatly simplifies the process of keeping OO repositories in sync.



On the **Runtime Environment** tab on the **Administration** tab, you can override the following configuration item types in individual OO installations:

- Remote Action Services (RAS)
- System accounts
- System properties

For information on adding a RAS reference, or creating a system account or system property, see the *Guide to Authoring Operations Orchestration Flows* (Studio_AuthorsGuide.pdf)

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

Account | Manage Users | Manage Groups | Dashboard Chart Definitions | System Configuration | Node Administration | Runtime Environment

Environment Overrides

Remote Action Service (RAS) Overrides | System Account Overrides | System Property Overrides

Refresh | Add | Edit | Remove

Name	Repository Value	Override Value
RAS_Operator_Path	https://BHOUSER1.americas.hpqcorp.net:9004/RAS/services/RCAgentService	http://ras2.domain.com

On the **Runtime Environment** tab, the **Environment Overrides** subtab shows the current overrides for all users. Changes to overrides take effect immediately on current and new flow runs.

When you change configuration items in the repository that have corresponding overrides in the runtime environment, these changes are synchronized with the OO environment. For example, if you delete a RAS definition, the override for that RAS is no longer valid. Later, if someone creates a RAS using the name of the deleted RAS, the override becomes active again and overrides the newly-created RAS. If you rename a configuration item, its corresponding override is also renamed and remains valid. If you need to remove an override, you can do it on the **Environment Overrides** subtab. Once you remove an override, the value of the configuration item in the OO repository is used.

Adding a RAS override

To add a RAS override

1. In Central, click the **Administration** tab.
2. On the Central **Administration** page, click the **Runtime Environment** tab.

The **Environment Overrides** page opens.

3. Click the **Remote Action Service (RAS) Overrides** tab.

Remote Action Service (RAS) Overrides tab

4. Click the **Add** button.

The **Add Remote Action Service (RAS) Overrides** dialog appears.

Add Remote Action Service (RAS) Overrides

Override Information

Name

UUID

Repository Value

Override Value

Override Information

Proxy None HTTP Proxy HP Gateway

5. In the **Name** box, select the RAS you want to override.
6. In the **Override Value** box, type the URL of the RAS you want to use as the override.
7. In the **Override Information** area, select the type of proxy to use for the override.

The **Proxy** options are **None**, **HTTP Proxy**, and **HP Gateway**.

- If you select **HTTP Proxy**, the **Add Remote Action Service (RAS) Overrides** dialog expands as shown below.

The screenshot shows a dialog box titled "Override Information". At the top, there are three radio buttons: "Proxy", "None", "HTTP Proxy", and "HP Gateway". The "HTTP Proxy" radio button is selected. Below the radio buttons are five text input fields: "Host", "Port", "Username", "Password", and "Retype Password". A mouse cursor is visible at the bottom right of the dialog box.

- In the **Host** box, type the name of the HTTP Proxy host.
- In the **Port** box, type the port number of the HTTP Proxy host.
- In the **Username** box, type the username for the HTTP Proxy host.
- In the **Password** and **Retype Password** boxes, type the password for the username.


- If you select **HP Gateway**, the **Add Remote Action Service (RAS) Overrides** dialog expands as shown below.

The screenshot shows a dialog box titled "Override Information". At the top, there are three radio buttons: "Proxy", "None", "HTTP Proxy", and "HP Gateway". The "HP Gateway" radio button is selected. Below the radio buttons are three text input fields: "Host", "Port", and "Realm".

- In the **Host** box, type the name of the HP Gateway host.
- In the **Port** box, type the port number of the HP Gateway host.
- In the **Realm** box, type the name of the user account's realm.

8. Click **Submit**.

The override appears in the **Remote Action Service (RAS) Overrides** tab.




Name	Repository Value	Override Value
RAS_Operator_Path	https://ras.americas.corp.net:9004/RAS/services/RCAgentService	https://ras2.domain.com:9004/RAS/services/RCAgentService

Adding a system account override

To add a system account override


1. In Central, click the **Administration** tab.
2. On the Central **Administration** tab, click the **Runtime Environment** tab.
The **Environment Overrides** page opens.
3. Click the **System Account Overrides** tab.



Name	Repository Value	Override Value

4. Click the **Add** button.

The **Add System Account Overrides** dialog opens.



Add System Account Overrides

Override Information

Name: SysAdmin

UUID: 1e412ee5-b76d-4b7a-949d-1152b9901eef

Repository Value: Username: System Administrator Password: ****

Override Value: Username:

 Password:

 Retype Password:

Cancel Submit

5. In the **Name** box, select the system account you want to override.
6. In the **Override Value** area:
 - a. In the **Username** box, type the system account username.
 - b. In the **Password** and **Retype Password** boxes, type the password for the username.
7. Click **Submit**.

The override appears on the **System Account Overrides** tab.


Adding a system property override

To add a system property override

1. In Central, click the **Administration** tab.
2. On the Central **Administration** tab, click the **Runtime Environment** tab.

The **Environment Overrides** page opens.

3. Click the **System Property Overrides** tab.



4. Click the **Add** button.

The **Add System Property Overrides** dialog opens.

The dialog box is titled 'Add System Property Overrides'. It contains the following fields:

- Override Information**
- Name:** A dropdown menu with 'username' selected.
- UUID:** A text field containing '77c0df42-a52b-4031-a8fa-7afcde2d8a8c'.
- Repository Value:** A text field containing 'sysadmin'.
- Override Value:** A text field containing 'admin'.

At the bottom of the dialog are two buttons: 'Cancel' and 'Submit'.

5. In the **Name** box, select the system property you want to override.

6. In the **Override Value** box, type the system property override value.
7. Click the **Submit** button.

Editing environment overrides

To edit an environment override

1. In Central, click the **Administration** tab.
2. On the Central **Administration** tab, click the **Runtime Environment** tab.
The **Environment Overrides** page opens.
3. Click the **Remote Action Service (RAS) Overrides** tab, **System Account Overrides** tab, or the **System Property Overrides** tab.
4. Select the override you want to edit.
5. Click the **Edit** button.
 - If you are on the **Remote Action Service (RAS) Overrides** tab, the **Edit Remote Action Service (RAS) Overrides** dialog opens.

Edit Remote Action Service (RAS) Overrides

Override Information

Name RAS_Operator_Path

UUID 5233030c-af46-432b-a682-b326ca6bf2ae

Repository Value https://BHOUSER1.americas.hpqcorp.net:9004/RAS/services/RCAgentService

Override Value

Override Information

Proxy None HTTP Proxy HP Gateway

- i. Edit those of the following fields that you want to change:
 - **Override Value.**
 - If **HTTP Proxy** is selected, you can edit the **Host**, **Port**, **Username**, and **Retype Password** fields.
 - If **HP Gateway** is selected, you can edit the **Host**, **Port**, and **Realm** fields.
 - ii. Click **Submit**.
- If you are on the **System Account Overrides** tab, the **Edit System Account Overrides**

dialog opens.

The screenshot shows a dialog box titled "Edit System Account Overrides". It contains the following information:

- Override Information**
- Name: SysAdmin
- UUID: 1e412ee5-b76d-4b7a-949d-1152b9901eef
- Repository Value: Username: System Administrator Password: *****
- Override Value: Username:
- Password:
- Retype Password:

At the bottom right, there are two buttons: "Cancel" and "Submit".

- i. Edit the following fields to specify the system account you want to be your environment override:
 - o **Username**
 - o **Password** and **Retype Password**
 - ii. Click **Submit**.
- If you are on the **System Property Overrides** tab, the **Edit System Property Overrides** dialog opens.

The screenshot shows a dialog box titled "Edit System Property Overrides". It contains the following information:

- Override Information**
- Name: username
- UUID: 77c0df42-a52b-4031-a8fa-7afcd2d8a8c
- Repository Value: sysadmin
- Override Value:

At the bottom right, there are two buttons: "Cancel" and "Submit".

- i. Edit the **Override Value** field to specify the value you want to use for the system override.
- ii. Click **Submit**.

Deleting environment overrides

To delete an environment override

1. In Central, click the **Administration** tab.
2. On the Central **Administration** tab, click the **Runtime Environment** tab.
The **Environment Overrides** page opens.
3. Click the **Remote Action Service (RAS) Overrides** tab, **System Account Overrides** tab, or the **System Property Overrides** tab.
4. Select the override you want to delete.
5. Click the **Remove** button.
The **Confirm Remove** dialog appears.
6. Click **OK**.

Best practices for runtime environment overrides

Overrides for RASes, system accounts, and system properties take effect immediately, so it's important to take proper care to ensure these changes don't have adverse effects on the respective Central installation.

The following best practices are recommended when making changes to the OO environment:

1. Make sure that all nodes are in the OFFLINE state and allow any existing flow runs to finish before you make changes to the environment. This ensures that flows don't pick up override changes in the middle of flow runs and produce inconsistent results. For example, adding a RAS override while a flow is running may result in part of the flow running on the original RAS and the rest of the flow running on the overridden RAS.

Note: If you make changes to the environment while flows runs are paused, the paused runs will use the latest environment state when they are resumed. However, override changes to system properties will not be used as the system properties are only added to the Global context of a run when the run is started.

2. Before you create, edit, or delete overrides, assess the impact of these changes by checking to see if any other items in the repository use the configuration item you want to override. Overrides apply to all flows that reference the overridden configuration item.

Other system configurations

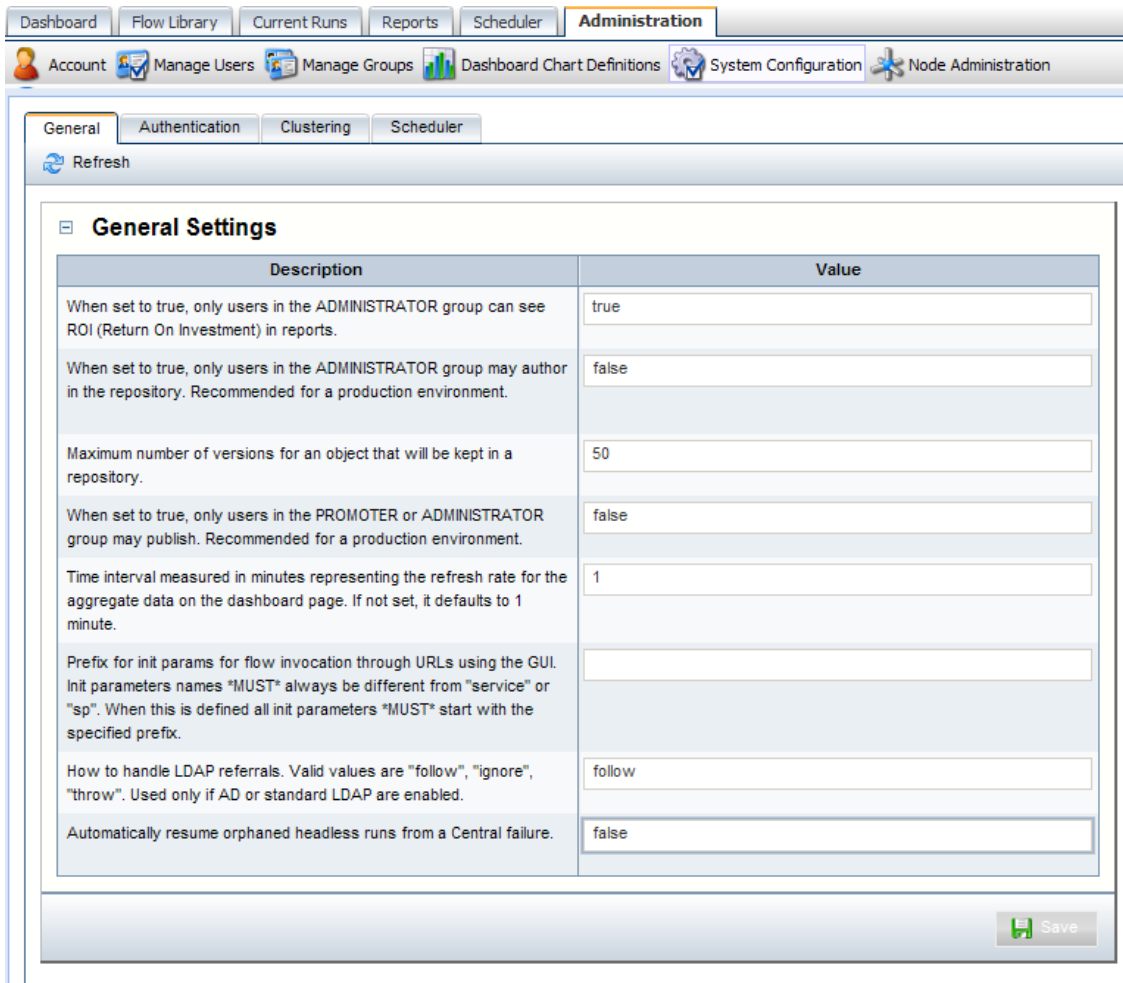
If you are a member of a group that has the `MANAGE_CONF` capability, you can customize several other aspects of the OO system:

- Whether Return on Investment (ROI) reporting is visible
- Whether only administrators can author in the Central repository
- The maximum number of versions of an object are kept in the repository
- Whether only members of `PROMOTER` or `ADMINISTRATOR` group can publish
- How frequently the Dashboard charts refresh
- For using a URL to start flow runs in the Central web page application, the prefix, if any, that should be required to start flow input names.

For information on specifying a prefix that must be used for all flow input names used in URLs, see the section below, [Specifying a required prefix for init params when using a URL to start a flow run in Central](#). Additional information is in the *OO Software Development Kit Guide* ([SDKGuide.pdf](#)).

- How LDAP referrals from the primary LDAP server to other LDAP servers are handled
- Whether to automatically resume headless runs that were interrupted by a Central server's failure
- How to enable load balancer management for a Central cluster

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



Enabling the visibility of ROI reporting

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

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

However, you can make ROI visible to all users.

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

To enable or disable ROI reporting visibility

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab.
3. On the **Administration** page toolbar, click the **System Configuration** tab.
4. On the **System Configuration** tab, click the **General** tab.
5. To allow all users to see ROI reporting, in the **When set to true, only users in the**

ADMINISTRATOR group can see ROI row, type **false** in the **Value** box.

OR

To allow only users in the ADMINISTRATOR group to see ROI reporting, in the **When set to true, only users in the ADMINISTRATOR group can see ROI** row, type **true** in the **Value** box. (This is the default setting.)

6. Click **Save**.
7. Restart the OO Central service (RSCentral).

Restricting authoring to administrative users

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

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

To restrict authoring to administrative users

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab.
3. On the **Administration** tab toolbar, click the **System Configuration** tab.
4. On the **System Configuration** tab, click the **General** tab.
5. To restrict the authoring capability to members of the ADMINISTRATOR group, in the **When set to true, only users in the ADMINISTRATOR group may author in the repository** row, type **true** in the **Value** box. The default value is **false**.

Note: If you have upgraded to OO version 9.00 from an earlier version, this value will default to true unless the default was changed to false in the previous version.

OR

To allow other users who are members of groups that have the AUTHOR capability, leave **false** in the **Value** box.

6. Click **Save**.
7. Restart the OO Central service (RSCentral).

Setting how many versions of an object are stored

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

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

To specify how many versions of an OO object are stored

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

Restricting who can publish

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

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

To restrict the ability to publish

1. Log on to Central with an account that has OO administrative permissions.
2. Click the **Administration** tab.
3. On the **Administration** tab toolbar, click the **System Configuration** tab.
4. On the **System Configuration** tab, click the **General** tab.
5. In the **When set to true, only users in the PROMOTER or ADMINISTRATOR group may publish** row, in the **Value** box, type **true**.

The default value is false.

Note: If you have upgraded to OO version 9.0 from an earlier version, this value will default to true unless the default was changed to false in the previous version.

6. Click **Save**.
7. Restart the OO Central service (RSCentral).

Changing the Dashboard charts refresh rate

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

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

To change the rate at which Dashboard charts refresh their data

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

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

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

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

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

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

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

Specifying how Central manages LDAP referrals

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

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

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

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

To specify how Central authentication manages LDAP referrals

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

Enabling Central to resume runs that were interrupted when Central failed

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

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

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

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

3. On the **Administration** tab toolbar, click the **System Configurations** tab.
4. On the **System Configuration** tab, click the **General** tab.
5. In the **Value** box for the **Automatically resume orphaned headless runs** setting, replace **false** with **true**.
6. Click **Save**.
7. Restart the OO Central service (RSCentral).

Enabling load balancer management for a Central cluster

If your deployment uses a load balancer and you want to synchronize OO state changes with it, you can enable load balancer management on the **Clustering** tab of the **System Configuration** page.


You must also create a flow that, each time a Central node's state changes (to ONLINE or OFFLINE), runs, updating the load-balancer's state.

For information on creating this flow, see [Integrating a load balancer with node management](#).

Note: To enable load balancer management for a Central cluster, your account must have the `MANAGE_CONF` capability.

To enable load balancer management for a Central cluster

1. In Central, click the **Administration** tab.
2. On the **Administration** tab toolbar, click the **System Configuration** tab.
3. On the **System Configuration** tab, click the **Clustering** tab.



The screenshot shows the 'System Configuration' page with the 'Clustering' tab selected. The 'Clustering Settings' section is expanded, displaying a table with the following settings:

Description	Value
The name of the cluster this Central belongs to.	CENTRAL_CLUSTER
Is Load Balancer management for the Central Cluster enabled.	false
The Load Balancer management flow for the Central Cluster (UUID or path).	/Library/Integrations/Hewlett-Packard/Operations Orchestration/Cluster/Samples/HP I

A 'Save' button is visible at the bottom right of the settings area.

4. For the **The name of the cluster this Central belongs to** setting, you specify the cluster name.

The cluster name you specify here is strictly for your own tracking purposes. It has no effect on which cluster the Central is actually a member of.
5. For the **Is Load Balancer management for Central Cluster enabled** setting, in the **Value** box, type **true**.

6. In the **The Load Balancer management flow for the Central Cluster** row, in the **Value** box, type either the path or the UUID of the flow that you created to manage the load-balancer.
7. Click **Save**.

Note: When you make and save any of the above changes, the following message appears above the **Clustering** tab. This message informs you that you must restart the OO Central service before the new settings will take effect.

 Changes Saved. Restart Operations Orchestration to use new Settings. 

Troubleshooting

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

You can safely proceed past this warning.

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

I was sent back to the login page

Your login may have timed out. Log in again.

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

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

I cannot change or create a schedule for a flow

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

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

Central fails to create a schedule for a flow

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

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

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

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

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

