

**Opsware Process Automation System
Central User Assistance**



Process Automation System Central

Opsware™ Process Automation System, Version 2.2

Users' Guide

For Process Automation System Central users:
Ops flow users and administrators

Copyright © 2000-2007 Opsware Inc. All Rights Reserved.

Opsware Inc. Unpublished Confidential Information. NOT for Redistribution. All Rights Reserved.

Opsware is protected by U.S. Patent Nos. 6,658,426, 6,751,702, 6,816,897, 6,763,361 and patents pending.

Opsware, SAS Web Client, Model Repository, Data Access Engine, Web Services Data Access Engine, Software Repository, Command Engine, Opsware Agent, Model Repository Multimaster Component, and Code Deployment & Rollback are trademarks and service marks of Opsware Inc. All other marks mentioned in this document are the property of their respective owners.

Additional proprietary information about third party and open source materials can be found at <http://www.opsware.com/support/sas65tpos.pdf>.

Where to Find Help, Tutorials, and More

The Process Automation System (PAS) documentation set is made up of the following:

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

This Help system and Guide

Help for PAS Central (reproduced in the PDF *Opware Process Automation System Central Users' Guide*, PAS_Central_UsersGuide.pdf) provides an introduction to Central and detailed procedures that you will use to create Ops flows.

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

- Introduction to PAS
This section is for all users; it gives an overview of PAS and its concepts.
- Using PAS Central
This section is for IT staff who run Ops flows
- Administering PAS
This section covers administrative tasks.

- Viewing Ops flow Reports and Audit Trails
This section is for IT managers who want to study metrics and reports on Ops flows and runs.

Quick View: PAS Central

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 Ops flows in the Central Library. An Ops 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 PAS charts and reports.

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

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

A scenario

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

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

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

We'll look at both these goals in turn.

PAS Central Web application: The Initial Page

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

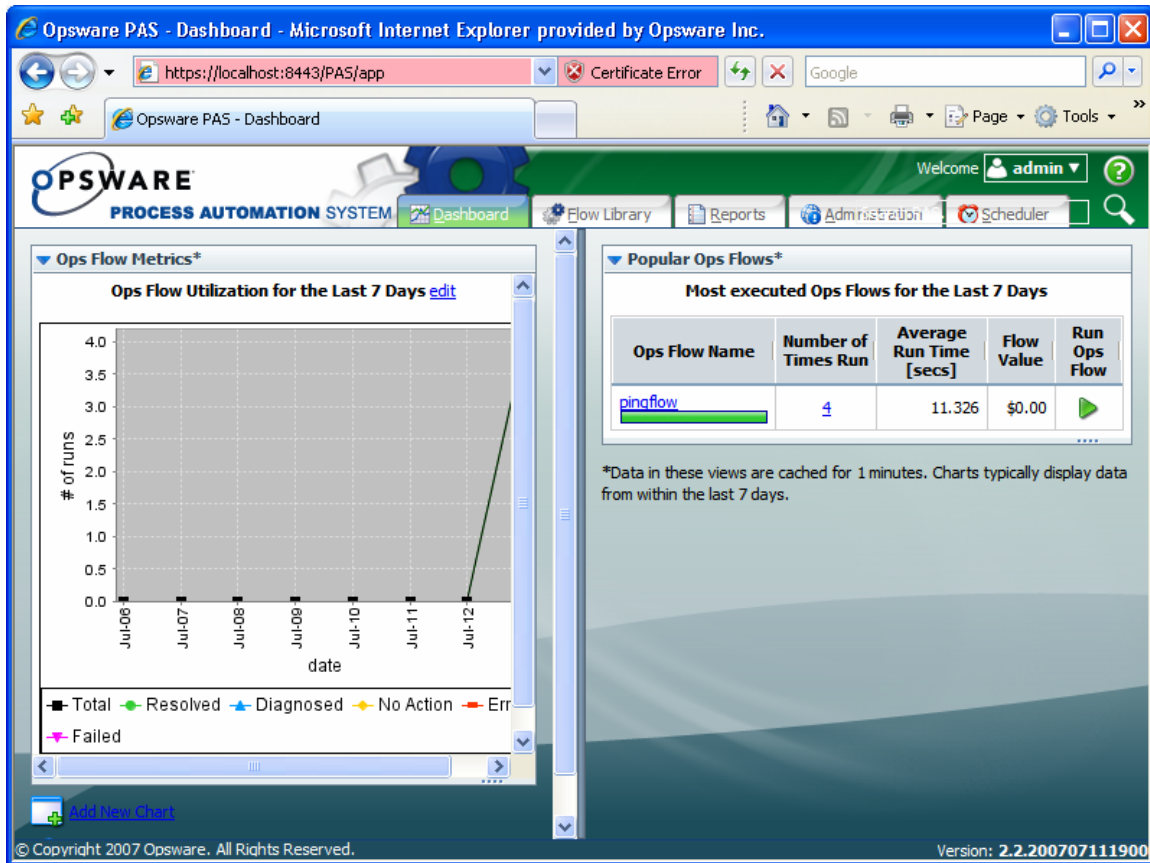


Figure 1 - Central Web application Dashboard

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

Flow Library tab: Choosing and running 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 Ops flow you needed to resolve each alert.

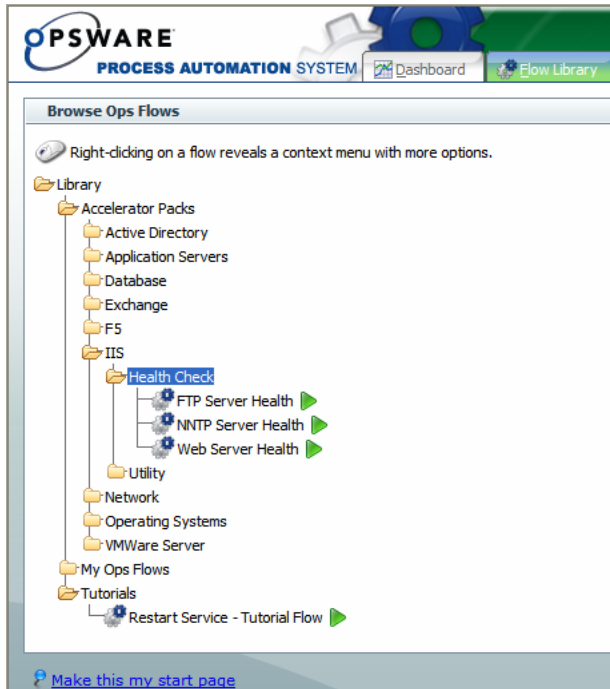


Figure 2 - Library with folder expanded

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

If you can't find the flow you need, try typing part of the name in the **Search** box in the upper-right corner of Central. You can run the flow directly from the results when they appear.



Figure 3 - Search box

When you have found your flow, you can run it by clicking either the right-pointing arrow (▶) or the name of the flow.

- If you click the arrow, the flow starts and runs to its end, pausing only for user prompts that you need to respond to, to provide the flow with information that it needs.
- If you click the flow name, you have the choice of running the flow step by step or running it to its end.

Tip: If you want to return to the Library without running the flow that you are previewing, just click the **Flow Library** tab again.

If you want to run the flow repeatedly, you can create a schedule for it.

Scheduling Ops flows

Suppose you need to regularly check whether a number of servers are online, you can schedule an Ops flow (say, "Connectivity Test") to start automatically at regular intervals that you define. Each schedule that you create can specify a different

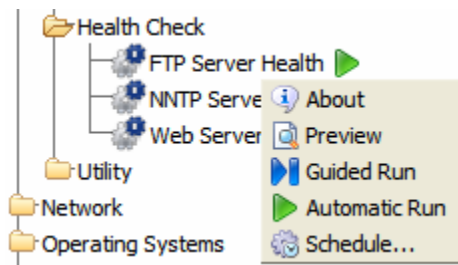
server's IP address for the flow to check. Creating schedules like this saves creating multiple flows to do the same thing, and saves you the work of starting each run individually.

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

For example, suppose that in Connectivity Test, the flow variable **host** stores the IP address of the server whose online status Server Status Flow should check. On the **Inputs** tab of the box in which you create the schedule, you would supply the IP address of the server you're interested in. For each subsequent schedule that you create for this flow, you would specify a different IP address.

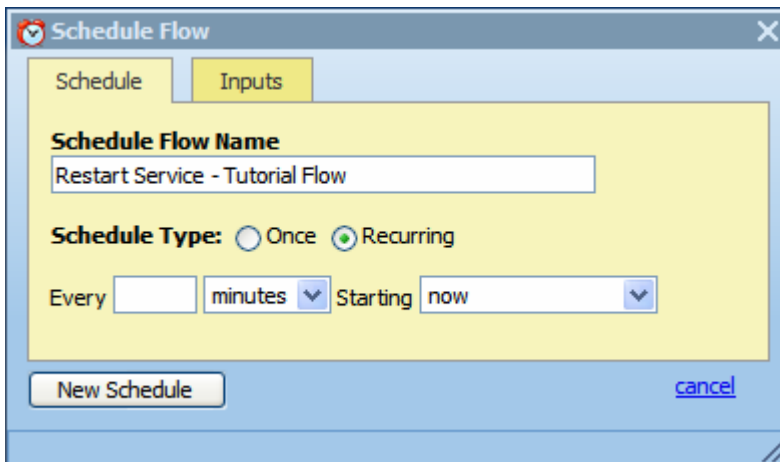
To create a schedule for a flow

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

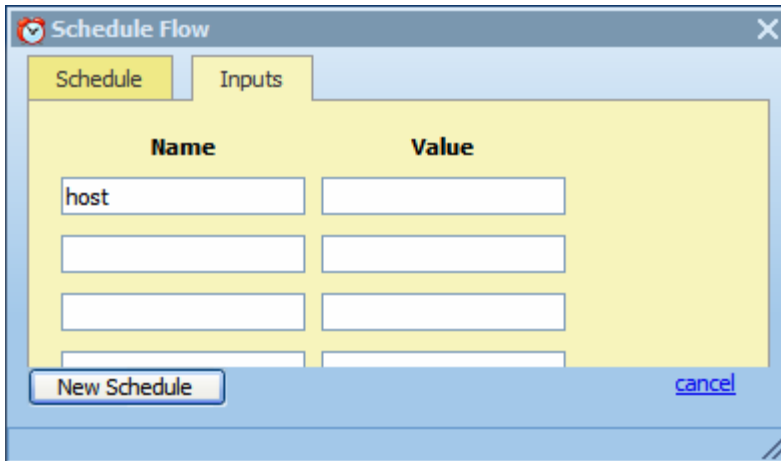


...and click **Schedule**.

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



3. On the **Inputs** tab...



...fill in the **Value** for all of the specified flow input names.

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

4. On either tab, click **New Schedule** to turn this specification into a real schedule. Once a schedule is created, you can edit it on the **Scheduler** tab.

Viewing and editing flow schedules

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

Scheduled Ops Flow	Enabled	Controls	Previous run result	Previous run time	Next run time	Delete
FTP Server Health /Library/Accelerator Pads/IIS/Health Check/FTP Server Health					Fri Jun 01 00:00:00 PDT 2007	
Restart Service - Tutorial Flow /Library/Tutorials/Restart Service - Tutorial Flow			failed to run schedule	Thu May 31 11:27:19 PDT 2007	Thu May 31 11:29:21 PDT 2007	

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

Figure 4 - Schedules on the Scheduler tab

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

Edit Schedule	Starting Date/Time	Ending Date/Time	Recurrence	Enabled	Controls	Previous run time	Next run time	Parameters	Delete
	Fri Jun 01 00:00:00 PDT 2007		every 30 minutes				Fri Jun 01 00:00:00 PDT 2007		
	Fri Jun 01 00:00:00 PDT 2007		every 30 minutes				Fri Jun 01 00:00:00 PDT 2007		
	Fri Jun 01 00:00:00 PDT 2007		every 30 minutes				Fri Jun 01 00:00:00 PDT 2007		
	Fri Jun 01 00:00:00 PDT 2007		every 30 minutes				Fri Jun 01 00:00:00 PDT 2007		

Figure 5 - Schedules for a single flow

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

Configuring Scheduler settings

You can control several aspects of how the Scheduler (or installations of the Scheduler) operate by specifying settings in the **Scheduler Settings** area of the **Administration** tab. To see the **Administration** tab, you must be logged in to Central with an account is a member of the ADMINISTRATOR group.

To specify Scheduler settings

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

The area looks like this:

Scheduler Settings	
Description	Value
The account that is used to run scheduled flows. It has to be an internal account that has HEADLESS_FLOWS capability.	admin
Log files directory (path is absolute). When scheduler is clustered, it should point to a network share that all nodes can read and write.	C:\Program Files\Opware\PAS\Scheduler\Vogs
Maximum log file size. Use KB,MB,GB	10MB
How many log files are retained	4
Log entry pattern	%-5p %d{MM/dd/yyyy HH:mm:ss} - %m%n
The size of most recent logging sent to the UI; if the log file is bigger than this size, only this amount from the end of the file is sent. Use KB, MB or GB	64KB
Number of flow input fields that the UI displays; increase this if a flow needs more than the default	10
If a flow xpath should be validated (slow process).	false
If PAS Scheduler(s) are clustered	true
The frequency (in ms) at which this instance "checks-in" with the other instances of the cluster. Affects the rate of detecting failed instances.	20000
Save Scheduler Changes	

Figure 6 – Configurations for the Scheduler

2. Complete the fields as needed.

As you complete the fields, note the following:

- You specified the account that is used to run scheduled flows when you installed the Scheduler. You can specify a different account here.

- The log settings are fairly self-explanatory. You should not tinker with the Log entry pattern unless you have a good understanding of the specifics of working with such patterns.
- The **number of flow input fields** is of interest. You cannot schedule a flow that has more flow inputs than the number specified in this setting, because the flow inputs that exceed what this setting allows cannot be given values by the scheduler, which means that the flow cannot run fully automatically, which ability is required by the Scheduler.

The last two settings address the possible clustering of multiple installations that run against a single PAS Central database, for load-balancing purposes:

- Always leave **If PAS Scheduler(s) are clustered** set to true. Enabling clustering does not affect performance, and even if you have a single installation of Scheduler, you may decide to install and cluster more in the future.
- The frequency with which the cluster of Schedulers check for failed Scheduler instances.

Quick Start: Starting a flow from outside Central

Rsflowinvoke.exe is a command-line executable that substitutes for a Web-browser call in starting an Ops flow. Rsflowinvoke.exe takes a URL as an argument, so you can run Rsflowinvoke from any machine as long as the machine you run it against, which is referenced in the URL that you pass as an argument, is accessible from the RAS or Central server that hosts the flow. This means that you can run Rsflowinvoke from a monitoring program such as MOM.

Rsflowinvoke.exe is useful when you want to start a flow from an external system, such as a monitoring application, that can use a command line to kick off a flow.

There are several different ways that you can start a flow with Rsflowinvoke.exe:

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

To start a flow using Rsflowinvoke

1. In a command line or in a script or batch file, use the following syntax for starting an Ops flow with Rsflowinvoke.exe:

```
RSFlowInvoke.exe <url> [-u -p -a -ep -rc -rw]
```

Where:

```
-u <username>
-p <password>
-ep <an existing encrypted password>
```

For the `-cp` option, with which you create an encrypted password, see below.

```
-a <authType>[Basic, Digest]
-rc <retry count> default=15 max=30
-rw <retry wait (seconds)> default=5
```

Example:

```
RSFlowInvoke.exe https://localhost:8443/PAS/services/http/list -u  
ofadmin -p password
```

2. To create an encrypted password, type the following:

```
Rsflowinvoke.exe -cp
```

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

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

To register or unregister in GAC

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

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

Where:

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

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

```
RSFlowInvoke.exe -s
```

Quick View: Learning more from Ops flows

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

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

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

Collating data on Dashboard reporting charts

You'll recall that we posed several questions that you might have after PAS has been running for a while. Charts that are available on the Dashboard tab can tell you:

- Which alerts are showing up most for each application and server
Note: PAS uses the ITIL term Configuration Item (or CI) to refer to server, applications and other items in your operations.
- Which server or application generated the most alerts?
The **Alerts per Configuration Item** chart answers both those questions.
- What actions have been taken on each application and server?

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

- Which flows have run the most often, and what were their outcomes?
Consider the **Outcomes per Flow** chart.
- Which flows were run to resolve errors, and how many times did the flows run?
Open the **All Alerts of Severity=Error Resolved by Ops Flows** chart.

You can bring these charts up on the **Dashboard** tab.

To bring up a chart with the current data

1. Click **Add New Chart**.
2. Select a **report** from the drop-down list and click **view**.

Tip: Among the domain terms that the charts record, here are meanings for the following:

Configuration Item

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

A Web server:

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

Your company home page:

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

Categories

The groups to which Ops flow authors assign flows. Charting categories enables you to view performance of these classes of flows. See for 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 Ops flows that produce the following chart. This chart shows you the outcomes, whose colors align with the flow return steps whose outcomes they represent (Diagnosed, No Action Taken, Resolved, and Failure).

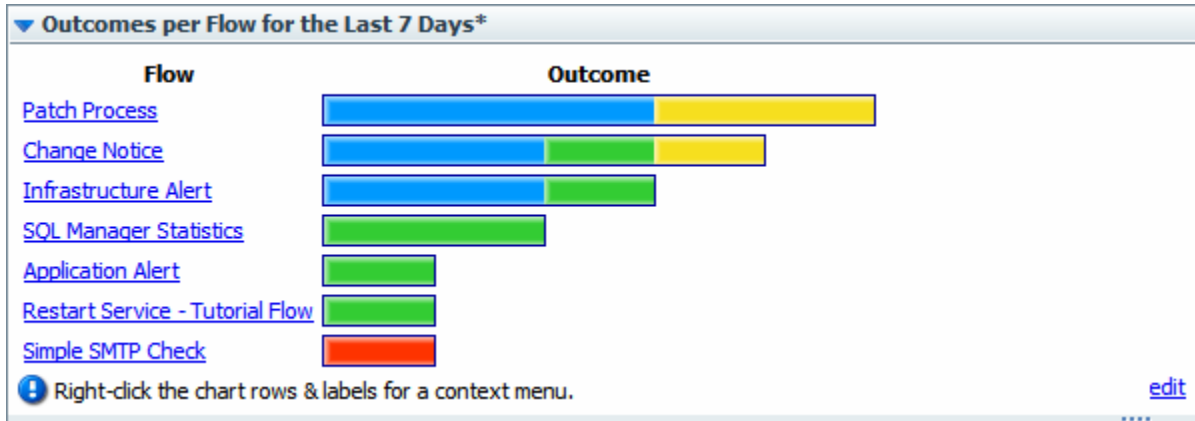


Figure 7 - Sample Dashboard chart

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

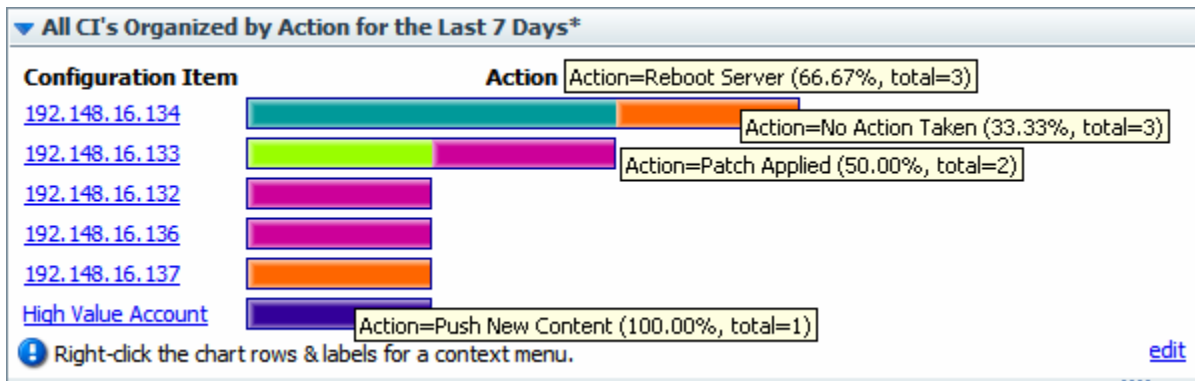


Figure 8 - Sample Dashboard chart with composite of bar labels

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

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

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

- Server 192.148.14.152 was rebooted by one flow and had no action taken by another.

- The Web application High Value Account had its content updated.
We can go further by drilling down into individual bars.

Learning more from the charts

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

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

Configuration Item	Action
192.148.16.134	

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

To learn more about data items in a chart

1. Right-click the appropriate bar segment for the data you’re interested in, then click **Drill Down**.
(**Tip:** You can also right-click the label at the left of the chart to drill down on all of the bars at once.)

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

The following box appears.

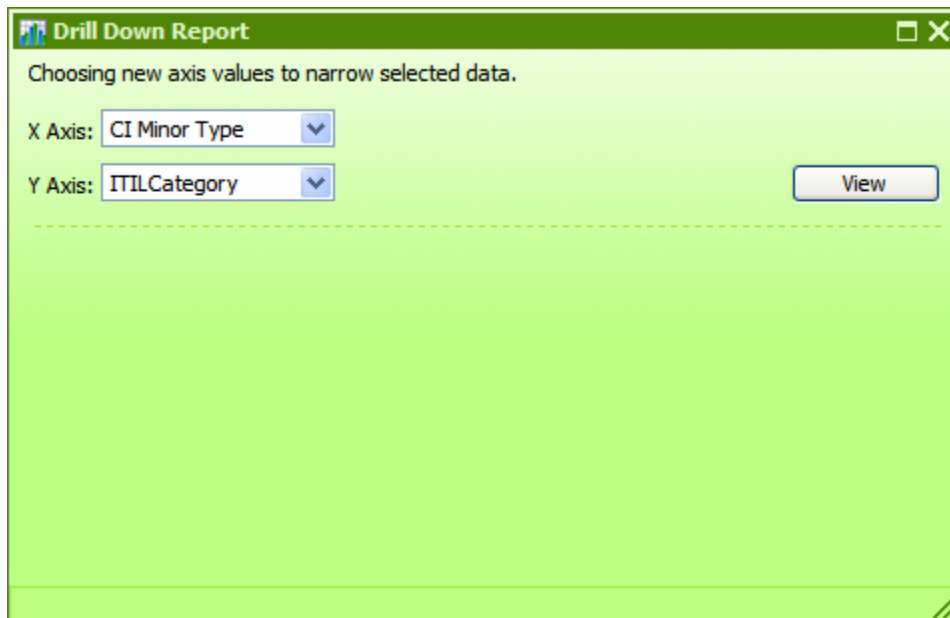


Figure 9 - Creating a drill-down report

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

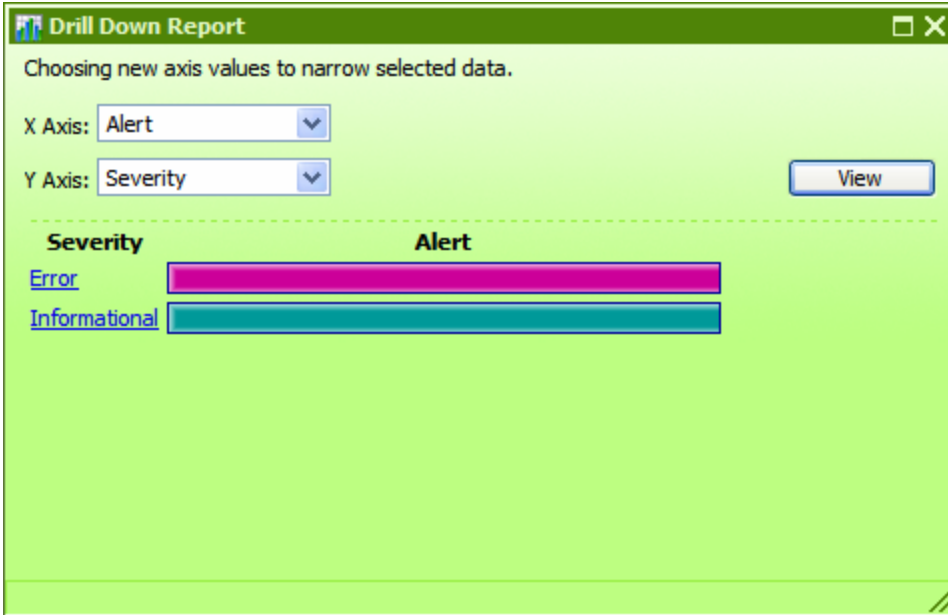


Figure 10 - The drill-down report that we created

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

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

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

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

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

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

To explore run-history reports, see [Run Reports tab: Examining run histories to see 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.

Making your own chart and changing existing charts

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

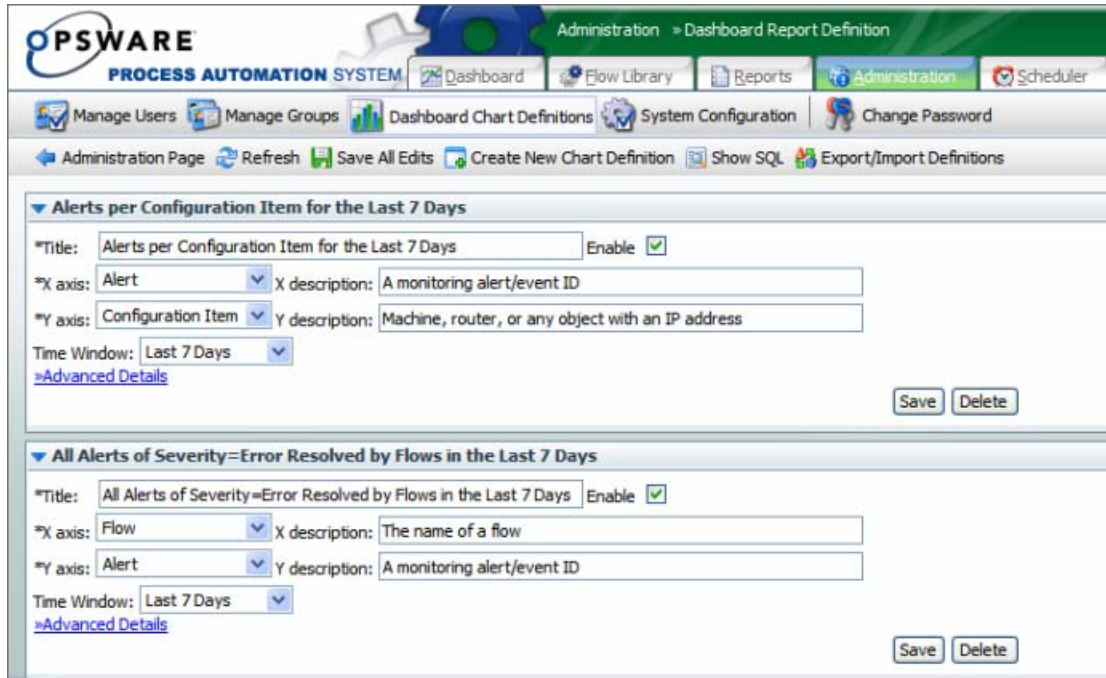
- Which applications and servers had fatal errors?
- How many alerts of various severities were there?
- How many alerts there are of each kind of severity (Informational, Warning, Error, Critical, Fatal)
- How many alerts of Fatal severity were there for each server and application?

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

To create a Dashboard chart

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

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



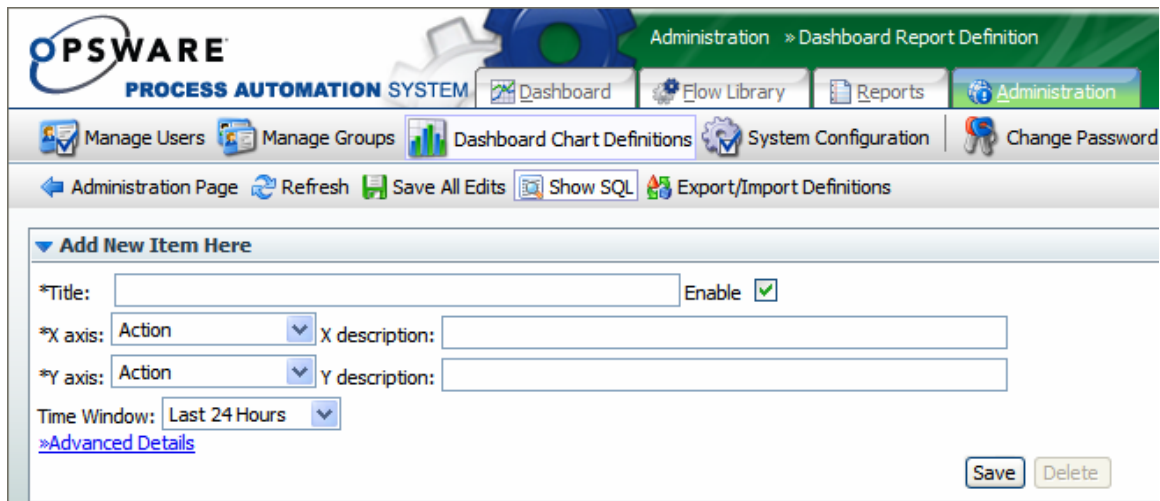
The screenshot shows the 'Administration > Dashboard Report Definition' page. The navigation bar includes 'Dashboard', 'Flow Library', 'Reports', 'Administration', and 'Scheduler'. The main menu has 'Manage Users', 'Manage Groups', 'Dashboard Chart Definitions', 'System Configuration', and 'Change Password'. The toolbar contains 'Administration Page', 'Refresh', 'Save All Edits', 'Create New Chart Definition', 'Show SQL', and 'Export/Import Definitions'. Two chart definitions are listed:

- Alerts per Configuration Item for the Last 7 Days**: Title: Alerts per Configuration Item for the Last 7 Days, Enable: . X axis: Alert, X description: A monitoring alert/event ID. Y axis: Configuration Item, Y description: Machine, router, or any object with an IP address. Time Window: Last 7 Days. Buttons: Save, Delete.
- All Alerts of Severity=Error Resolved by Flows in the Last 7 Days**: Title: All Alerts of Severity=Error Resolved by Flows in the Last 7 Days, Enable: . X axis: Flow, X description: The name of a flow. Y axis: Alert, Y description: A monitoring alert/event ID. Time Window: Last 7 Days. Buttons: Save, Delete.

Figure 11 - Administration tab, Dashboard chart definitions

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

A new chart definition box appears.



The screenshot shows the 'Administration > Dashboard Report Definition' page with the 'Add New Item Here' form. The navigation bar and main menu are the same as in Figure 11. The toolbar includes 'Administration Page', 'Refresh', 'Save All Edits', 'Show SQL', and 'Export/Import Definitions'. The form fields are:

- Title**: [Empty text box], Enable:
- X axis**: Action, X description: [Empty text box]
- Y axis**: Action, Y description: [Empty text box]
- Time Window**: Last 24 Hours
- Buttons: Save, Delete

Figure 12 - Creating a new Dashboard chart

3. Type a title for the new chart.
4. In the **X axis** drop-down list, select what you want to chart on the horizontal axis, and then type a description.

- In the **Y axis** drop-down list, select what you want to chart on the vertical axis, and then type a description.
- 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.

▼ Add New Item Here

*Title: Enable

*X axis: X description:

*Y axis: Y description:

Time Window:

[»Advanced Details](#)

Comments: Optional comments

Top X: X Threshold: Top Y: Auto-group

Restrict X axis values:

Group all other values in one segment:

Name for all other values segment:

Additional Constraints...

Domain Term Name: Domain Term Value:

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

- 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.
- 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.
- 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.

10. Under **Additional Restraints**, from the **Domain Term Name** drop-down, select the domain term charted on the Y axis, and then type a value in the **Domain Term Value** box.

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

11. On the **Administration** tab, click **Dashboard Chart Definitions**.
12. Scroll down to the chart you want to change.
13. In the box that defines the chart you want to change, make any desired changes, and then click **Save All Edits**.

Notes:

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

Examining run histories to see what happened and why

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

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

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

Parameters for Reporting Level: All Ops Flow Types

The screenshot shows a 'Report Definition' window with the following sections:

- Time Window:** Includes a radio button for 'Relative' (selected) with a dropdown menu set to 'Last 24 Hours', and a radio button for 'Absolute' with 'Start' and 'End' date pickers and 'Time' fields set to '00:00'.
- Filters:** Includes a 'Subtree' dropdown menu set to '<all folders in repository>', an 'Executed by User(s)' text field, and a 'Result' section with checkboxes for 'All', 'Diagnosed', 'Resolved', 'No Action Taken', 'Error', and 'Failed'.
- Matching Step Inputs:** Includes 'Input Name:' and 'Input Value:' text fields with a '+1 More' button.
- Matching Domain Values:** Includes 'Domain Term:' and 'Domain Term Value:' text fields with a '+1 More' button.

At the bottom of the window are 'Search' and 'New Search' buttons.

Figure 13 - Defining reports of run histories

To define which run histories are examined

1. Specify the time window over which you want to see run histories.
For a fixed window, pick start and end dates by clicking the calendars. Let's click the calendars and pick 09/04/06 for the **Start** date and 09/06/06 for the **End** date.
If you specify the time, do so in 24-hour format.
2. Under **Filters**, in the **Subtree** drop-down list, pick the Library path that contains the Ops flow whose runs you want to see.
The drop-down list contains only the subtree paths that contain Ops flows that have been run.
In our example, we'll pick **/Library/Accelerator Packs/Windows Management**.
Although for this example we'll omit specifying users, you could do so, separating the user names with commas.
3. Select the **Result** check box or boxes that you want the report to include.
4. Under **Matching Step Inputs**, add any input names and values you want to use to further limit the runs included in the report.
5. To specify another step input name and value, you click **+ 1 More**, then repeat the preceding step (#4), specifying the second server's IP address or name.
6. If there is a domain term you want to use to further define the search, follow the same steps to specify domain term name/value pairs that you did for step inputs.

7. Under **Report Columns** on the left, de-select the columns you're not interested in.

8. Click **Search**.

The search results appear.

▼ Report Data										
Found 7 results at 06/01/07 11:24:12										
Ops Flow Name	Number of Runs ▼	Avg. Number of Steps	Avg. Repair Time [secs]	Most Recent Run	Percent Resolved	Percent Diagnosed	Percent Error	Percent No Action	Percent Failed	Flow Value
Patch Process	5	6.00	.196	05/31/07 13:12:54	0.00%	60.00%	0.00%	40.00%	0.00%	\$83.83
Change Notice	4	6.00	.343	05/31/07 13:09:59	25.00%	50.00%	0.00%	25.00%	0.00%	\$0.00
Infrastructure Alert	3	7.00	.434	05/31/07 13:11:14	33.33%	66.67%	0.00%	0.00%	0.00%	\$53.47
SQL Manager Statistics	2	7.00	34.873	05/31/07 12:21:18	100.00%	0.00%	0.00%	0.00%	0.00%	\$0.00
Application Alert	1	7.00	.862	05/31/07 13:07:24	100.00%	0.00%	0.00%	0.00%	0.00%	\$16.17
Restart Service - Tutorial Flow	1	90.00	10.285	05/31/07 11:45:15	100.00%	0.00%	0.00%	0.00%	0.00%	\$2.15
Simple SMTP Check	1	4.00	62.334	05/31/07 11:48:47	0.00%	0.00%	100.00%	0.00%	0.00%	\$0.00

Figure 14 - Search results

From here you can drill down to a single flow and a single run.

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

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

▼ Report Data							
Found 5 results at 06/01/07 11:35:08							
Flow Id	Start Time ▼	User Id	Time To Run [secs]	Run Result	Number of Steps	Flow Path	Run Value
50	05/31/07 13:12:54	admin	.220	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	\$16.43
49	05/31/07 13:12:37	admin	.211	No Action Taken	6	/Library/My Ops Flows/Dashboard/Patch Process	\$17.27
48	05/31/07 13:12:19	admin	.160	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	\$16.43
47	05/31/07 13:12:03	admin	.181	Diagnosed	6	/Library/My Ops Flows/Dashboard/Patch Process	\$16.43
46	05/31/07 13:11:41	admin	.211	No Action Taken	6	/Library/My Ops Flows/Dashboard/Patch Process	\$17.27

Figure 15 - Run history of one flow

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

We're interested in the run that Resolved, so we'll click that Flow ID (**4**).




Figure 16 - Run report for a single flow run

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

PAS Central: Detailed Reference

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

You access and run Ops flows from Central to:

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

Before you continue reading, make sure that you have familiarized yourself with the basics of PAS and Ops flows in Chapter 1, "Introduction and road map."

Central

Central provides a graphical user interface for:

- Finding and running Ops flows.
- Creating reports and viewing information on Ops flow runs.

Starting Central

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

To start Central

1. Start your machine's Web browser.

2. 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**.
A message appears, warning you that the site is not trusted. However, it is safe to proceed.
4. Click **Yes**.
5. When the Central **Login** page appears, log in with your user name and password. Central opens and you're ready to locate, run, and view information on Ops flows.

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

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

Ops flow Metrics area

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

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




Figure 17 - Ops Flow Metrics area

To customize the Ops flow Metrics graph

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

The Ops 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.

Ops Flow Utilization for the Last 7 Days

Ops Flow Utilization

Last 7 Days

Go

[close](#)

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 Ops flows

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

▼ Popular Ops Flows*

Most executed Ops Flows for the Last 7 Days

Ops Flow Name	Number of Times Run	Average Run Time [secs]	Flow Value	Run Ops Flow
Patch Process	5	0.196	\$83.83	
Change Notice	4	0.343	\$0.00	
Infrastructure Alert	3	0.434	\$53.47	
SQL Manager Statistics	2	34.873	\$0.00	
Application Alert	1	0.862	\$16.17	
Simple SMTP Check	1	62.334	\$0.00	
Restart Service - Tutorial Flow	1	10.285	\$2.15	

Add New Chart

Figure 18 - Popular Ops Flows

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

To open a Dashboard chart

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

Navigating in Central

Central varies according to whether you are finding or running an Ops flow or generating a report or metrics. However, you can always navigate with the **Dashboard**, **Ops Flows**, **Run Reports**, **Administration**, and **Help** tabs.



Figure 19 - Navigation tabs

Finding an Ops flow

You can find and open an Ops flow by either:

- Browsing the Library
- Searching

You can also open a flow by typing or pasting its URL directly into your Web browser.

Browsing Ops flows in the Library

To browse the Library for an Ops flow

1. Click the **Flow Library** tab.
The Ops flow **Library** opens.

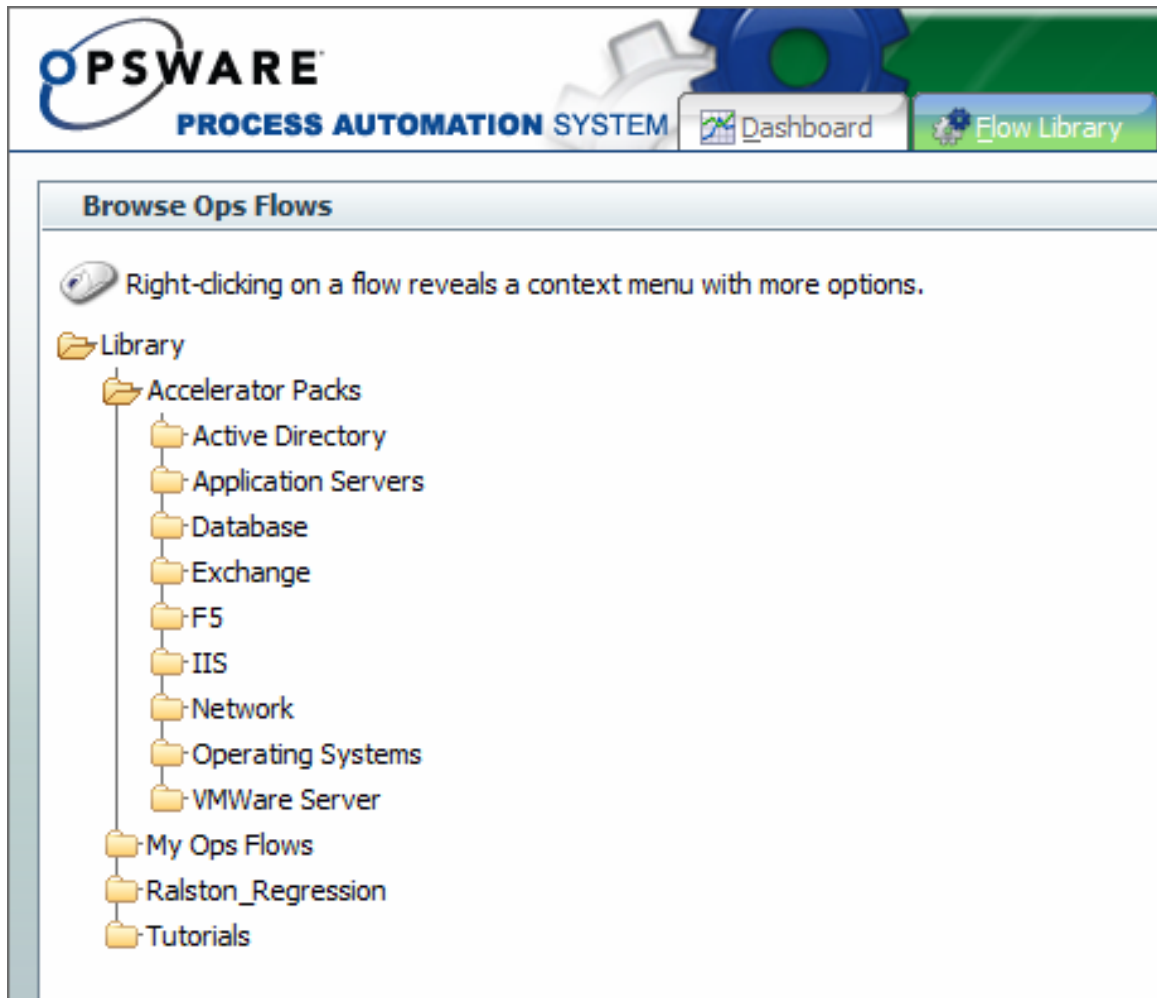


Figure 20 - The Flow Library

2. To find an Ops flow, open the **Library** folder and navigate through the folder tree to the flow.


The  icon represents an Ops flow.

Tip: To see short descriptions of what each flow does, click **Show Short Descriptions**. (When the descriptions are displayed, the command changes to **Hide Short Descriptions**.)

There are two ways to run an Ops flow:

- Step by step—you click to carry out each step and respond to any user prompts.
- To completion—you only respond to user prompts; the flow completes each step on its own.

For more on running a flow, see [Running Ops flows](#).

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

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

4. To run the flow, do one of the following:

- Click the green arrow to the right of the name.
This loads the flow and runs it to completion.
- Click the flow name
OR

In the “i” balloon information box, click **run Ops flow**.

This loads a preview of the flow and allows you to choose between running the flow to completion and running it step by step.

Searching for an Ops flow

You can search for an Ops flow from anywhere in Central, using the **Search PAS** box.



Figure 21 - Search PAS box

To search for an Ops flow

1. In the **Search PAS** box, type any of the following:
 - The name of the flow
 - Keywords
 - A word or phrase within the flow description
 - A flow category
2. Press ENTER on your keyboard.
Central shows the results of the search.

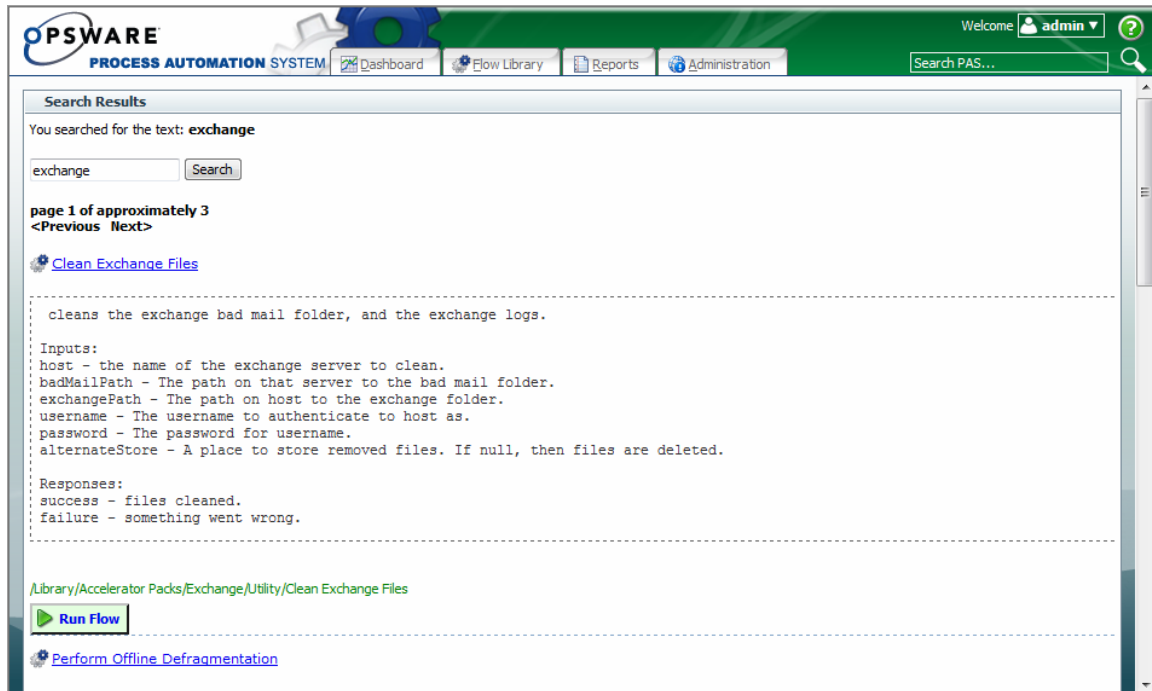


Figure 22 - Search results

3. Scroll down to the flow that you want to run.
4. To run the Ops flow to completion, click **Run Flow** beside the green arrow.

OR

To load a preview of the flow and have the choice of running it to completion or step by step, click the flow name.

Ops flow previews

The Ops flow preview page, which appears when you click the Ops flow name elsewhere in Central, contains the flow diagram and information about it. When you run the Ops flow step by step, the diagram illustrates the current progress of the flow.

The Advanced Details area, when you click the down arrow,

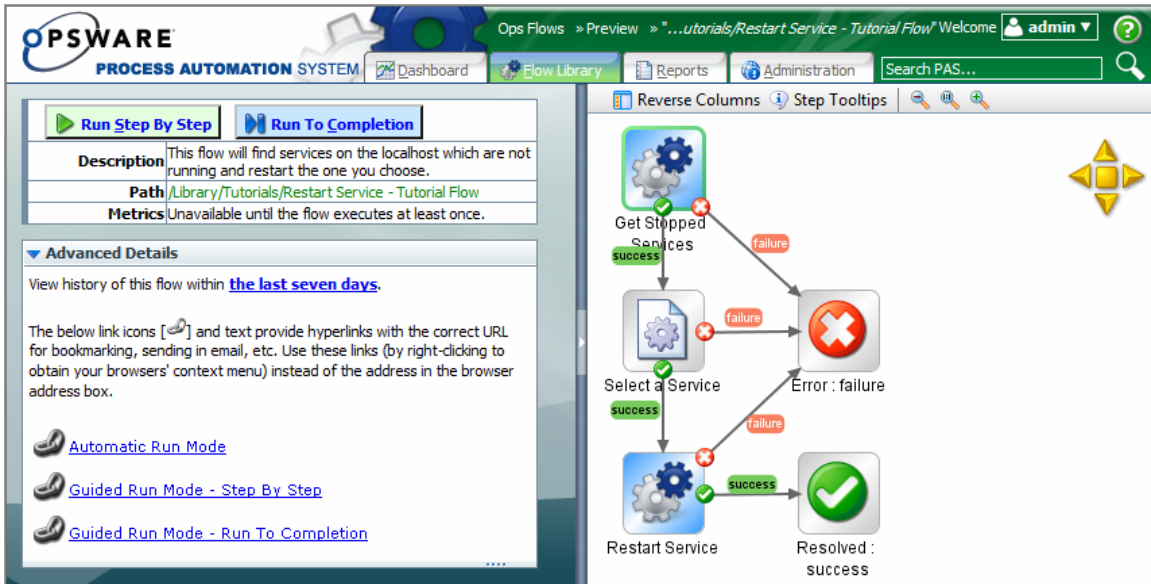







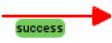
Figure 23 - Previewing a flow

To show more of the Ops flow diagram

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

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

Table 1. Flow diagram symbols and their meaning

Symbol	Name	Meaning and comments
	Start step	The entry point for an Ops flow
	Diagnosed return step	The step that ends an Ops flow when a problem has been diagnosed
	Error return step	The step that ends an Ops flow when an error has occurred that prevents the Ops flow from continuing
	Resolved return step	The step that ends an Ops flow when a problem has been resolved
	No action taken return step	The step that ends an Ops 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.



Ops flow move buttons

Click the directional buttons to move the diagram flow.



Reverse Columns button

Click this to reverse the position of the flow diagram and the flow details pane.

To move the Ops flow diagram on the page

1. Click the movement arrows in the top-left of the Ops flow diagram.
OR
Drag the diagram.
2. To return the Ops flow to its original position, click the center movement button.

To reverse the panels on the preview page

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

Advanced Details

Before you start a flow run, the **Advanced Details** area contains useful links; after you start a flow, the area contains step information that is updated with the completion of each step in the flow.

To expand or contract the Advanced Details area

- Click the down arrow on the left.

Before you start a run, the **Advanced Details** area appears as follows:

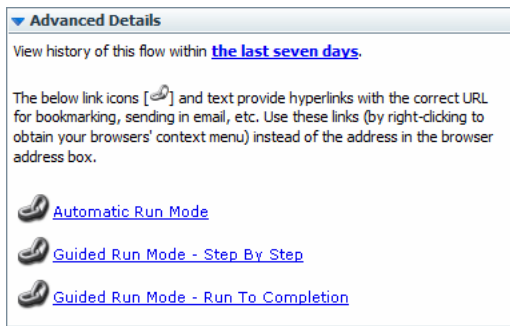


Figure 24 - Advanced Details, preview

The links in the area open a run history of the Ops flow for the last seven days and links that start the run in the various run modes. For more information on the run modes, see [Running Ops flows](#).

Running Ops flows

There are three ways to run a flow: to completion, in either simple or advanced mode, or (in advanced mode) step by step.

- Running a flow to completion means that the flow runs to one of its end points (return steps), stopping only for the user to provide information required by a user prompt.

- Running step by step means that the user must initiate completion of each step, in addition to responding to user prompts.

The two modes are distinguished by how much information is provided to the user.

- Simple mode provides only the information in the **Results Summary** area after completion of the flow. Simple mode always runs the flow to completion
- Advance mode provides more information in the **Advanced Details** area.

To run a flow to completion

- On the **Flow Library** tab, by the flow name, click the green arrow or **Run Ops flow**.

To run a flow step by step

1. On the flow's preview page, click **Run Step by Step**.

As the flow runs, the results of each step are displayed under **Completed Steps**, as in the following example. If the Ops flow has passed data to the context, this is reported in the **Advanced** area.

As you run each step, the **Advanced** area below the Ops flow diagram is updated with information about the step.

2. Click **Next** to advance the flow to each next step.

If you need to supply input, a dialog box appears, with drop-down lists for you to choose from or text boxes for you type responses into, depending on what information the flow needs.

3. Supply the input and then click **Continue**.

OR

If you don't have the necessary information available, click **Cancel** and then either interrupt the run to obtain the information and resume it later or hand the run off to another user who has the information.

For information on interrupting a run, see [Interrupting Ops flow runs](#).

For information on handing off a run, see [Handing off Ops flow runs](#).

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

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

5. As each step is completed, the outcome of the step appears in the **Description** column under **Results Summary**, as shown in the following illustration:





▼ Results Summary		
Step	Response	Message
Get Stopped Services		Retrieved a list of services which are currently stopped.
Select a Service		Selected Adobe LM Service to restart.
Restart Service		Restarted service Adobe LM Service
Resolved : success		Return step - Restart Service - Tutorial Flow

Figure 25 - Results Summary of a flow run

Running subflows

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

To step into and out of a subflow

1. When the step highlight moves to the step that contains the subflow, click **Step Into**.
2. Complete the steps of the subflow using the same procedure as you do for completing the steps of the parent flow.
3. To run the subflow to completion and return to the steps of the parent flow, click **Return**.

OR

Click **Step out**.

Starting an Ops flow from a URL

To create a linked URL that can start a flow run

1. Click the **Flow Library** tab, navigate to the flow, and click the flow name to open the preview of the flow.
2. Under **Advanced Details**, right-click **Automatic Run Mode**, and then click **Copy Link Location**.
A copy is created of the URL that launches the flow.
3. To send the URL to another Central user, paste the URL into a message.
4. If the flow has any required inputs, modify the URL by adding name-value pairs that define values for all the inputs.

For information on defining a prefix for input-value pairs, see [Defining a prefix for inputs in URLs that launch flows](#).

5. To create a name-value pair for each input, append to the URL a string with the following syntax:

```
&<initparamprefix><inputname>=<inputvalue>
```

Where

```
<initparamprefix>
```

is the prefix that you or your administrator specified in the General Settings area of the Central Administration tab's System Configuration subtab.

```
<inputname>
```

is the name of the input.

```
<inputvalue>
```

is the value to use for the input in the flow run that the URL initiates.

Note that the input name does not have to be enclosed by `{` and `}`.

For example, if your flow included a ping operation and so needed a target, then to ping the server "ram," you might set a flow input named target to the value "ram". If your init param prefix were "__", you would do so with the following:

```
$_target=ram
```

Interrupting Ops flow runs

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

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

To interrupt an Ops flow run

- Click **Interrupt**.

Central brings up the **Run Administration** tab.

Resuming a run

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

To resume an Ops flow run

1. Click the **Run Administration** tab.

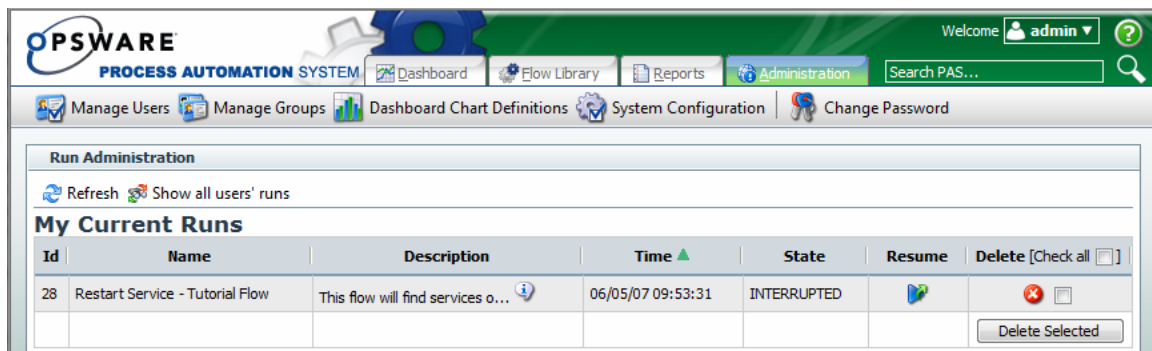


Figure 26 - Run Administration

2. In the **Resume** column, click the resume icon () for the appropriate flow run.


Canceling a run

To cancel an Ops flow run

- Click **Cancel**.

Creating a link to a run

You might want bookmark an Ops flow run or create a shortcut to it, which you can place on another Web page, in a document, or in an email message.


To do so, you do not use the address in your Web browser address box, but the hyperlinked text that is identified with a link icon ()

There are three link icons, one for each kind of run that you want to execute:


- Automatic Run Mode
- Guided Run Mode – Step by Step

- Guided Run Mode –Run to Completion

To create a link to an Ops flow from an external source

1. On the **Flow Library** tab, right-click the hyperlinked text beside the appropriate  icon.
2. In the context menu that appears, click **Copy Shortcut**.
3. In the external source from which you want to access the run, paste the shortcut that you copied.

To bookmark the Ops flow link in your Web browser

1. On the **Flow Library** tab, right-click the hyperlinked text beside the appropriate  icon.
2. In the context menu that appears, click your browser's command that bookmarks the link, and then complete your browser's process for doing so.

Handing off Ops flow runs

You might need to hand off an Ops flow in which:

- A step requires information that someone else has.
- A transition is gated (requires access permissions that your account does not possess).

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

To hand off an Ops flow

1. After starting a step by step run of the flow that you want to hand off, click **Hand Off** in the left-hand panel.
The run is paused and the state of the run is changed to **Handed Off**.
A new email message appears, with the URL of the Ops flow included in the body of the message.
2. Address the message to the person to whom you're handing off the Ops flow, and send the message.

To resume an Ops flow run that has been handed off to you

1. Open the email message that contains the URL of the Ops flow and click the URL.
A new copy of the Web browser opens, on the **Administration** tab, where the handed-off flow is added to the list of incomplete flows under **Run Administration**.

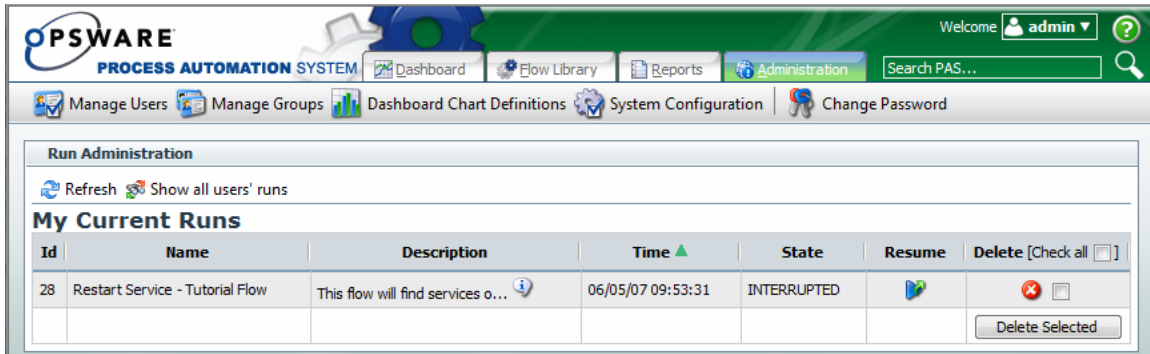



Figure 27 - Run Administration

2. In the **Resume** column, click the resume icon () for the appropriate flow run.

Auditing and managing Ops flows

Audit information on Ops flows (individually or groups) and their runs can be particularly important in system diagnostics at several levels, for Central users, PAS administrators, and IT managers. For information on auditing and viewing reports on flow runs, see “Auditing and Viewing Ops flow Reports.”

While users of Central can interrupt or cancel, resume or restart, Ops flow runs that they own, administrative privileges are necessary to view all current Ops flow runs and pause, resume, or delete them. For information on these and other administrative tasks, see “Administrative tasks.”

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

Individual flow authors grant access permissions to their flows in studios. For information on assigning users or groups access permissions for various PAS 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 PAS for working with authentication by Active Directory, Lightweight Directory Access Protocol, or Kerberos providers and add PAS users, in [Getting 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 [Configuring log file settings](#) and [Configuring heap size for PAS server process](#).

Capabilities and access permissions

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

- A capability is the right to perform an action in PAS, such as the `MANAGE_USERS` and `MANAGE_GROUPS` capabilities. A PAS administrator (a user with these just-mentioned capabilities) assigns groups the capabilities that they need. For more information, see [Capabilities](#).
- Permissions are access rights to individual objects, such as individual folders, flows, operations, or system accounts. The four permissions are `READ`, `WRITE`, `EXECUTE`, and `LINK`, which flow authors grant to users or groups for individual objects. So:
 - To find and run an Ops flow in Central, users must have read and execute permissions for the flow. In Studio, authors must have read, write, link, and/or execute permissions for objects that they use to author flows. For instance:
 - To debug a flow, an author must have the execute permission for that flow.
 - A flow author must have the Link permission for any flow or operation from which he or she creates a step in a flow.
 - To change a system account, an author must have the Read and Write permissions for the system account.

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

Capabilities

Following are the capabilities that can be assigned in PAS.

Capability	Description
<code>MANAGE_USERS</code>	Create, delete, and modify internal users. Only holders of this capability can create internal PAS users.
<code>MANAGE_GROUPS</code>	Create, delete, and modify groups.
<code>MANAGE_CONF</code>	Manage configuration properties and dashboards.
<code>AUTHOR</code>	Start Studio.
<code>SCHEDULE</code>	Schedule Ops flows.
<code>VIEW_SCHEDULES</code>	View Ops flow schedules.
<code>MANAGE_RUNS</code>	View, delete, and reassign runs other than the user's.
<code>RUN_REPORTS</code>	Run reports and view metrics Dashboard pages.
<code>RUN_HEADLESS</code>	Start flows from the Web Service interface.

Permissions

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

Permissions for PAS objects

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

PAS objects and the permissions needed to work with them

Object	Action	Necessary permission(s)
Folder		
	View contents	Read, write
	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	Link to
System account		
	View account name	Read
	Change account password	Read, write

	Rename account	Read, write
	Use in flow or operation	Link to
	Use at runtime	Execute

For more information on the groups, capabilities, and permissions model of PAS security, see the *PAS Administration Guide*.

When you first deployed PAS, you mapped users to the PAS groups. Depending on how you accomplish the mapping, when you deploy the PAS 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 a PAS group. For information on mapping users to PAS group, see the installation and deployment guide, *Installing Opsware Process Automation System* (PAS_InstallGuide.pdf).

Getting external users into the Central system

Besides creating users within PAS (called *internal* users), you can map external users or groups to PAS groups. To map external users to PAS groups, you first specify that PAS authenticate external users with your organization's authentication provider (Active Directory, Lightweight Directory Access Protocol [LDAP], or Kerberos).

Using external authentication for Central users

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

To enable PAS for one or more authentication providers

1. To enable authentication provider, select the appropriate check box (**AD Enabled**, **LDAP Enabled**, or **Kerberos Enabled**).
2. Modify the configuration values for the authentication provider according to your organization's needs, according to one of the sections that follow this procedure:
 - [AD authentication settings](#)
 - [LDAP authentication settings](#)
 - [Kerberos authentication settings](#)
3. After configuring settings according to one of the sections listed in the preceding step, test the current settings for an authentication provider from within PAS by clicking the **Test** button for the authentication provider.

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

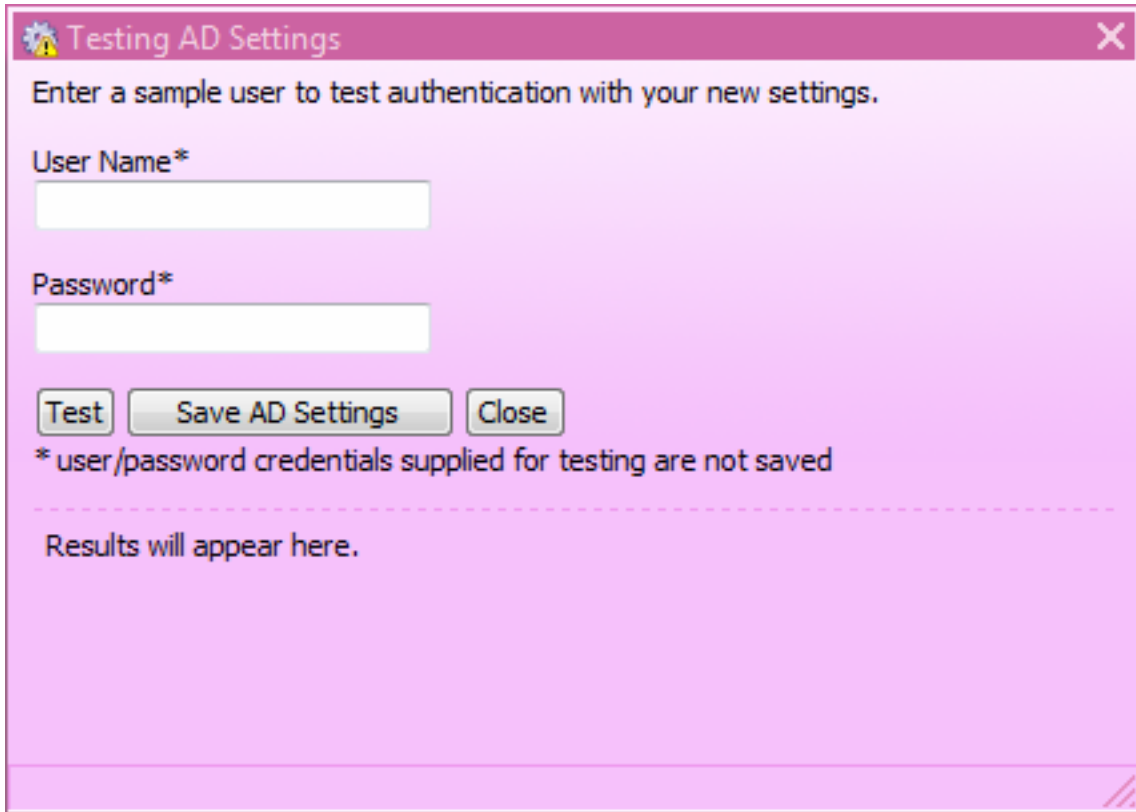


Figure 28 - Testing an authentication provider

4. Type the user name and password of an external account that is authenticated by the provider you're testing, and then click **Test**.
5. If the test fails, modify the settings and test again.
6. When the authentication provider settings test successfully, click the **Save AD Settings** button for the authentication provider.
7. When you see a message to restart the RSCentral service, do so.
8. After configuring the system,

For the following sections on configuring PAS using AD or LDAP authentication, suppose the following users are members of these groups:

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

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

- Its IP address is 192.111.5.102

- Its fully qualified name is mirage.ad

AD authentication settings

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

AD Authentication Settings AD Enabled

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

Figure 29 – AD Authentication Settings

To authenticate with Active Directory

1. Select the **AD Enabled** check box.
2. Supply the URL or IP address of the AD, with the following syntax:

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

where:

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

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

`LDAP://mirage.ad:200`

or

`LDAP://192.168.5.5:200`

Important: Machines ordinarily communicate with Active Directory using Lightweight Directory Access Protocol (LDAP, a clear-text protocol). To encrypt communications, you can set PAS 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 *PAS Administration Guide* for information on configuring your system to use the LDAPS protocol.

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

`LDAPS://mirage.ad:200`

or

LDAPS://192.168.5.5:200

3. In the **Value** box for the **List of LDAP contexts containing user groups**, type the contexts in which LDAP should search for your existing, external groups.

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

For example, suppose that:

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

The following setting makes these groups visible to PAS:

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

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

4. In the **Value** box for **LDAP filters that try to match the user groups**, type a filter to find the groups for users.

For instance, suppose the following:

- Tom Gage's "memberOf" entry has a value of "CN=Service Desk,CN=Users,DC=mirage,DC=com"
- Mary Grey's "memberOf" entry has a value of "CN=Network Specialist,CN=Users,DC=mirage,DC=com"
- Ed Stuart's "memberOf" entry has a value of "CN=Manager, OU=Staff, DC=mirage, DC=com"

Thus Tom's and Mary's LDAP entries are defined under the context CN=Users, while Ed's entry is defined under OU=Staff.

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

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

Be sure to:

- Type each instance of **member=CN={1}**, exactly as it appears above.
- Type the semicolon (;) separator between the filters that you type.

5. Leave **name** in the **Value** box for the **Attribute** setting.
6. To specify the contexts in which PAS should look for users, type the contexts in the **Value** box for the **List of LDAP contexts containing users** setting.

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

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

7. In the **Value** box for the **user domain** setting, type the domain where the users reside.

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

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


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

- In the **Value** box for the **default role**, specify which PAS group or individual is mapped to when a mapping is not specified.

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

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

AD Authentication Settings AD Enabled

Description	Value
Active Directory URL.	LDAP://mirage.com
List of LDAP contexts containing user groups. For example, if you have a "Groups" object containing groups of users, then the expression: ou=Groups,dc=mycompany,dc=com might be used. The list separator is a ";".	CN=Users,DC=mirage,DC=ad
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
Attribute of any group (returned from the group search), to use as group name.	name
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
The user domain (the authentication used is the NT style "domain\username").	mirage\\{0}
The default group an AD authenticated user gets when there is no group matching.	LEVEL_ONE

Save AD Settings Test Refresh

Figure 30 - AD authentication enabled

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

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

LDAP authentication settings

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

LDAP Authentication Settings LDAP Enabled

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

Save LDAP Settings Test Refresh

Figure 31 - Settings for configuring PAS with LDAP authentication To authenticate with LDAP

10. Select the **LDAP Enabled** check box.

11. Supply the URL or IP address of the top level of the LDAP server, with the following syntax:

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

where:

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

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

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

or

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

Important: To encrypt communications, you can set PAS to communicate with LDAP over SSL by specifying the LDAPS protocol. For information on configuring your system to use the LDAPS protocol, see the *PAS Administration Guide*.

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. In the **Value** box for the **List of LDAP contexts containing user groups**, type the contexts in which LDAP should search for your existing, external groups.

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

For example, suppose that:

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

The following setting makes these groups visible to PAS:

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

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

13. In the **Value** box for **LDAP filter that tries to match the user groups**, type a filter to find the groups for users.

For instance, suppose the following:

- Tom Gage's "memberOf" entry has a value of "CN=Service Desk,CN=Users,DC=mirage,DC=com"
- Mary Grey's "memberOf" entry has a value of "CN=Network Specialist,CN=Users,DC=mirage,DC=com"

- Ed Stuart's "memberOf" entry has a value of "CN=Manager, OU=Staff, DC=mirage, DC=com"

Thus Tom's and Mary's LDAP entries are defined under the context CN=Users, while Ed's entry is defined under OU=Staff.

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

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

This filter finds groups in either:

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

or

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

Be sure to:

- Type each instance of **member=CN={1}**, exactly as it appears above.
- If you type more than one filter, separate the filters with a semicolon (;).

14. Leave **name** in the **Value** box for the **Attribute** setting.

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

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

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

16. In the **Value** box for the **default group**, specify which PAS group or individual is mapped to when a mapping is not specified.

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

17. To save your settings, click **Save LDAP Settings**.

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

Kerberos authentication settings

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

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

Kerberos Authentication Settings Kerberos Enabled

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

Figure 32 - Settings for configuring PAS with Kerberos authentication
To authenticate with Kerberos

1. Select the **Kerberos Enabled** check box.
 2. In the **Value** box for the **Kerberos 5 configuration file**, type the name of the file.
 The location of the file should be within the PAS home directory, and the path should be relative to that directory.
 For instance, the example in the description to the left of the **Value** box describes the Kerberos configuration file as being in the \Central\conf subdirectory of the PAS home directory.
 3. In the **Value** box for the **KDC host**, type the IP address or fully qualified machine name of the Key Distribution Center (KDC), the authentication center for users.
 Use the following syntax:

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

 where:
 - **<KDC_host>** is the IP address or fully qualified name of the LDAP server.
 - **<port>** is the port number that the KDC host uses, if the KDC host is configured to use a non-default port.
 For instance, if the LDAP server is mirage.ad, its IP address is 192.168.5.5, and it uses port 200, the setting would be:

```
mirage.ad:200
```

 or

```
192.168.5.5:200
```
 4. In the **Value** box for the **Kerberos realm**, type the domain name of the realm.
 For instance, the domain might be **MIRAGE.AD**.
 5. In the **Value** box for the **default group**, specify which PAS group or individual is mapped to when a mapping is not specified.
 For instance, to assign unmapped groups or individuals only the capability to run flows, you would type **LEVEL_ONE**.
 6. To save your settings, click **Save Kerberos Settings**.
- If your IT organization also authenticates with Active Directory and/or LDAP, complete the procedures in [AD authentication settings](#) or [LDAP authentication settings](#).

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

Managing users

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

- When you create an internal user, you also create a password and assign the user to one or more PAS 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 PAS groups, or map the user's external (AD or LDAP) account to a PAS account.

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

Adding a user

To add an individual user

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

The **User Information** dialog appears.

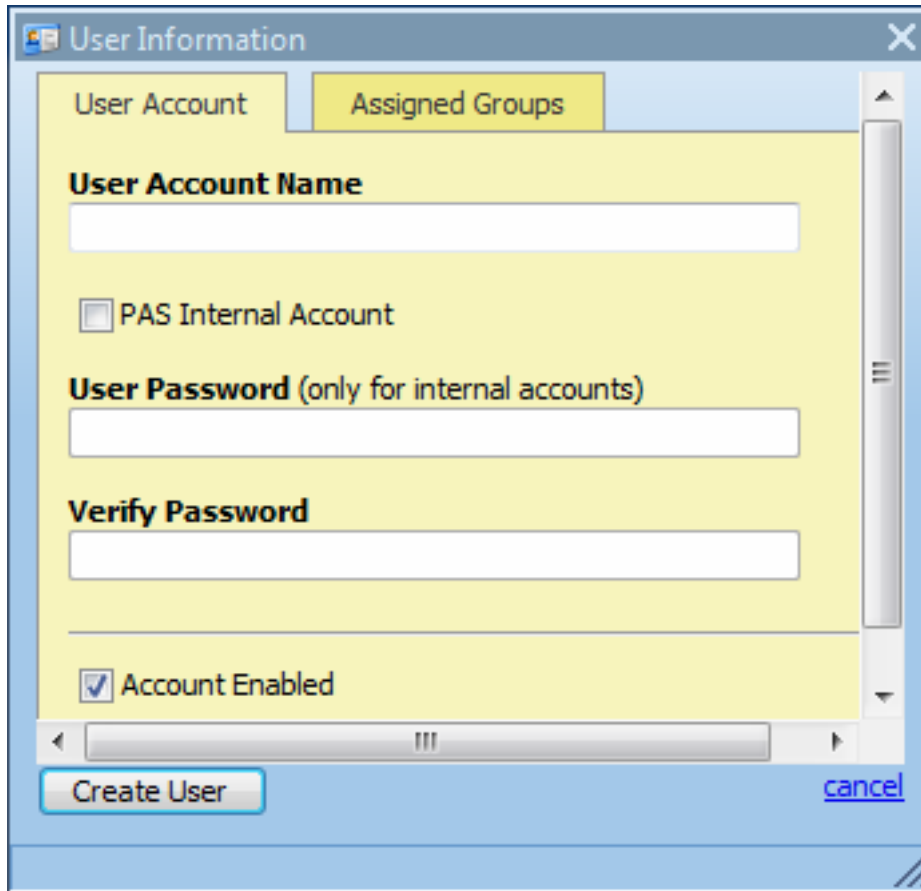


Figure 33 - User Information dialog box

2. Type the user account name.
3. If you are creating the account within PAS, select the **PAS Internal Account** check box, and then type and verify a password for the user to log in to PAS with.
Note that by default, the **Account Enabled** checkbox is selected.
4. To assign the new user to a group, click the **Assigned Groups** tab within the dialog and then select the check boxes for the groups whose capabilities the user should have. For information on the capabilities that are assigned to the various groups, see the *Administering Opsware Process Automation System* (PAS_AdminGuide.pdf).

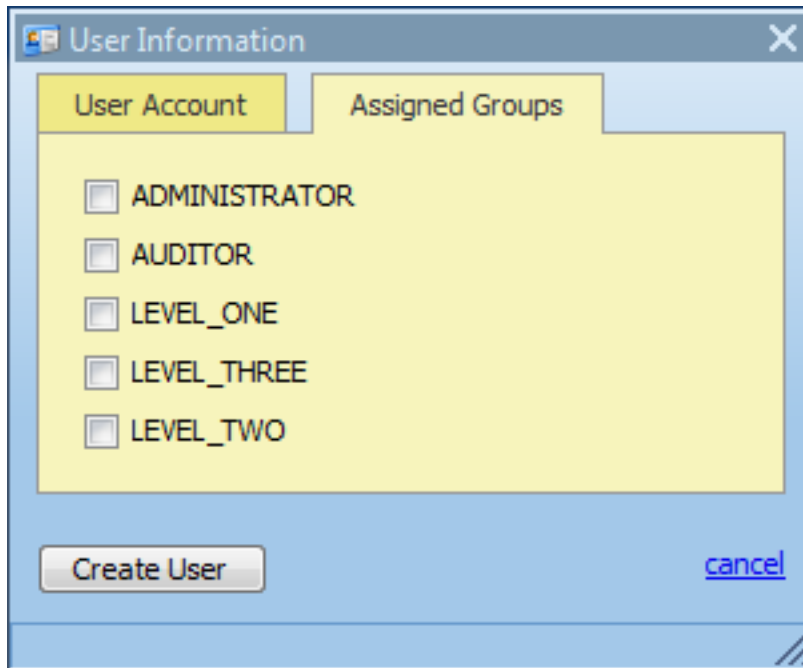



Figure 34 - Assigning groups to a user

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

Editing a user's account

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

To make changes to a user's account

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

Deleting a user

To delete a user account

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

Managing groups

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

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

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

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

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

Scenario

Suppose you want to map the following groups in your IT organization to existing PAS 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 PAS, 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 PAS groups and create rights for the PAS groups that make sense for these external groups, let's look at the groups that are created by default when you install Central:

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

The following are three special PAS groups:

- ADMINISTRATOR

The purpose of the ADMINISTRATOR group is to have one account that you can use to run and work in Central and Studio in case you temporarily lose the ability to log users in. This group possesses all capabilities and access permissions, so you cannot modify its capabilities or access permissions. You can, however, change the password.

Although you can add members to this group, keep in mind that it is an all-powerful group within PAS, 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 PAS automatically becomes a member of this group. The group doesn't have any capabilities, but does have access to certain PAS

objects, such as Accelerator Packs. As a result, the PAS 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.

Groups					
Administration Page Refresh Add New Group					
Edit	Group Name ▲	Description	Capabilities	External Groups Mapping	Delete [Check all <input type="checkbox"/>
	ADMINISTRATOR	Represents PAS administrators.	MANAGE_USERS, MANAGE_GROUPS, AUTHOR, SCHEDULE, MANAGE_RUNS, RUN_REPORTS, MANAGE_CONF, VIEW_SCHEDULES, HEADLESS_FLOWS		
	AUDITOR	Represents PAS 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 PAS level one users.	NONE		<input type="checkbox"/>
	LEVEL_THREE	Represents PAS level three users.	NONE		<input type="checkbox"/>
	LEVEL_TWO	Represents PAS level two users.	NONE		<input type="checkbox"/>

Delete Selected

Figure 35 - Manage Groups subtab

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


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

Mapping external groups to PAS groups

To add users of the following AD or LDAP groups to PAS groups as shown in this table, use the **Manage Groups** subtab on the **Administration** tab.

AD or LDAP group	PAS group
Service Desk	LEVEL_ONE
Network Specialist	LEVEL_THREE
Manager	ADMINISTRATOR

To map external groups to PAS groups

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

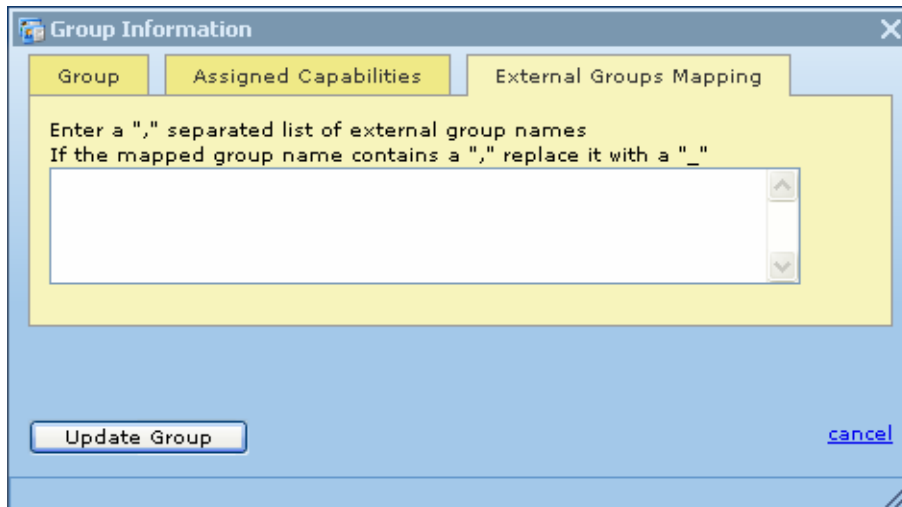


Figure 36 – External Groups Mapping tab

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

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

After making the above group assignments, the

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

Permissions are granted by Ops flow authors in Studio. The available permissions are read, write, execute, and link permissions to Ops flows and objects that are associated with them. For more information on granting permissions, see Help for Studio.

To add a PAS user

1. Click the **Administration** tab, then click the **Manage Users** subtab.
2. Click **Add New User**.
3. Complete the information in the **User Information** dialog box’s **User Account** tab.

To assign a user to a PAS group

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

Changing a group’s assigned capabilities and description

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

To change the capabilities assigned to a group

- In the **Group Information** dialog, click the **Assigned Capabilities** tab.

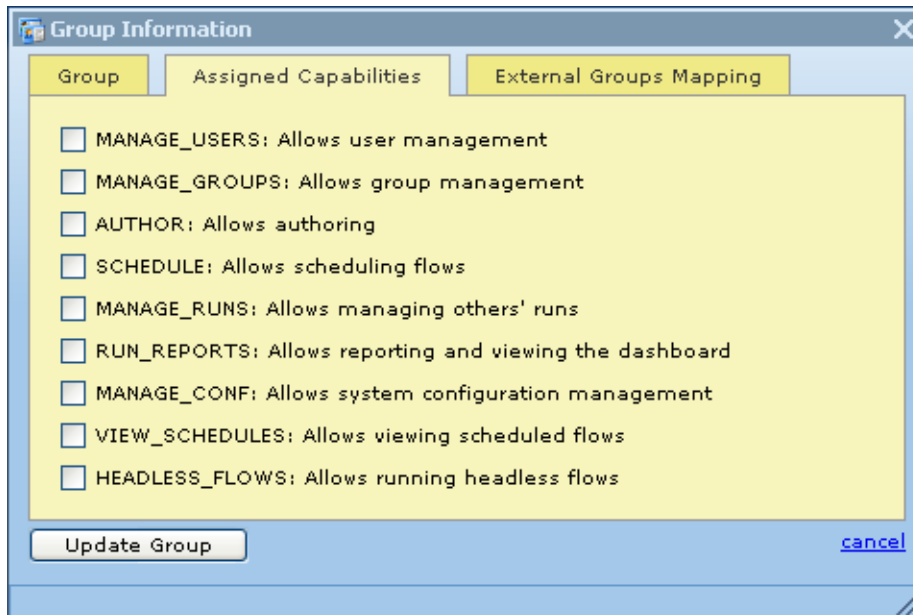


Figure 37 - Assigning capabilities for a group

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

Notes:

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

In addition, the following tasks must be executed outside the Central application. For more information on these tasks, see *Administering Opsware Process Automation System*.

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

Changing Scheduler settings

You can only change settings for the Scheduler settings if the Scheduler is installed.

Other configuration changes

On the **Administration** tab, you can specify that authors who do not have administrative privileges can directly connect to the public repository. However, by default, this is not permitted. It is strongly recommended that you leave this default setting as is, because allowing non-administrative authors to directly connect the the public repository decreases the security of PAS and your flows, both creating a security vulnerability and reducing the control over the changing of flows used in your production environment.

Managing Ops flow runs

On the **Administration** page, you can view current Ops flow runs and resume, delete, or reassign them.

Note: To open the **Administration** page, you must be a member of the PAS ADMINISTRATOR group. For adding user accounts to PAS groups, see [Managing groups](#).

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

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

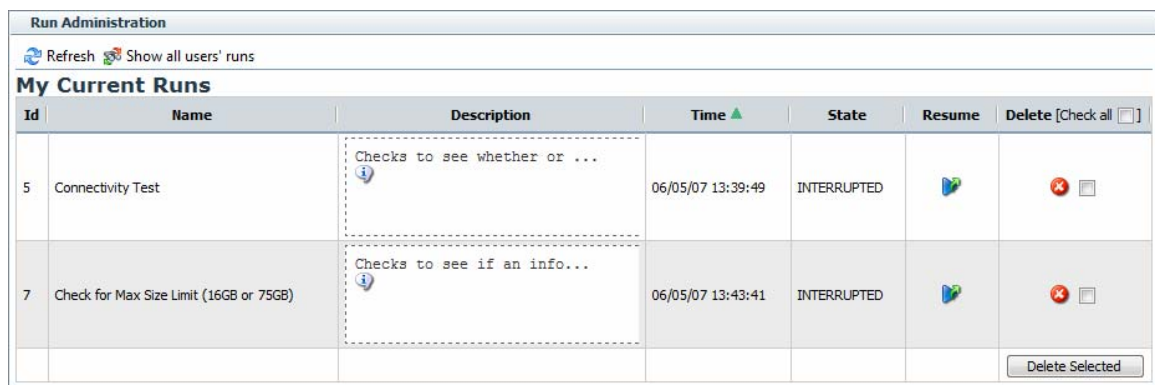
To see which Ops flows are currently running

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

OR

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

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



The screenshot shows the 'Run Administration' interface. At the top, there are 'Refresh' and 'Show all users' runs' buttons. Below is a section titled 'My Current Runs' containing a table with columns: Id, Name, Description, Time, State, Resume, and Delete. Two rows are visible, both with 'INTERRUPTED' state. The first row has Id 5 and Name 'Connectivity Test'. The second row has Id 7 and Name 'Check for Max Size Limit (16GB or 75GB)'. Each row has a 'Resume' button (green play icon) and a 'Delete' button (red X icon). A 'Delete Selected' button is located at the bottom right of the table.


Id	Name	Description	Time	State	Resume	Delete [Check all <input type="checkbox"/>
5	Connectivity Test	Checks to see whether or ...	06/05/07 13:39:49	INTERRUPTED		
7	Check for Max Size Limit (16GB or 75GB)	Checks to see if an info...	06/05/07 13:43:41	INTERRUPTED		

Figure 38 - Administering Runs

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

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

To resume a run

- In the **Resume** column, click the resume arrow ().

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


To delete a run

- Select the **Delete** check box, and then click **Delete Selected**.

To reassign a run

- Under **User**, type the user name of the person to whom you want to assign the run, and then click **Submit Changes**.

Important: If the user name is misspelled, the run is associated with the misspelled name and the run is not visible to the intended user. To correct this, retype the name with the correct spelling.

- Under **Resume**, click the resume arrow ().

For a **Warning** about the states for which it is safe to resume a run, see the procedure above, "To resume a run."

Other system configurations

There are several other PAS system configurations that you can change:

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

Enabling ROI reporting

By default, ROI reporting is enabled.

To enable or disable ROI reporting

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

Changing the Dashboard charts refresh rate

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

To change the rate at which Dashboard charts refresh their data

7. Log on to Central with an account that has PAS administrative permissions.
8. Click the **Administration** tab, and then the **System Configuration** tab.
9. To change how frequently Dashboard charts are updated with new data, in the **General Settings** area, in the **Time interval...refresh rate** row, in the **Value** box, type a whole number reflecting the number of minutes you want between updates.

Defining a prefix for inputs in URLs that launch flows

To start a flow run from a URL, the URL must include the flow's inputs, with values defined. The syntax for doing so is:

```
&<initparamprefix><inputname>=<inputvalue>
```

Where

`<initparamprefix>`

is the prefix specified in the General Settings area of the Central Administration tab's System Configuration subtab.

`<inputname>`

is the name of the input.

`<inputvalue>`

is the value to use for the input in the flow run that the URL initiates.

Use the following procedure to define the prefix.

To define a prefix for specifying name and value pairs for input parameters in URLs that start flow runs

10. Log in to Central with an account that has PAS administrative rights.
11. On the **Administration** tab, click the **System Configuration** subtab.
12. In the **General Settings** area, in the **Prefix for init params** row, in the **Value** box, type a prefix that you will use in the URL that you send to your recipient.
As noted in the **Description** box for this row, the prefix must not be "service" or "sp".
13. Click **Save General Settings**.

Specifying how Central manages LDAP referrals

When you have enabled Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) authentication, you can specify how referrals from what Central

authentication does when it encounters an LDAP referral from one server or namespace to another. You can specify that Central do one of the following:

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

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

To specify how Central authentication manages LDAP referrals

14. Log on to Central with an account that has PAS administrative permissions.
15. Click the **Administration** tab, and then the **System Configuration** tab.
16. On the **How to handle LDAP referrals** line, in the **Value** box, type either **follow**, **ignore**, or **throw**, depending on how you want Central authentication to respond.

Other Administration Configurations

On the **Administration** tab, you can specify that authors who do not have administrative privileges can directly connect to the public repository. However, by default, this is not permitted. It is strongly recommended that you leave this default setting as is, because allowing non-administrative authors to directly connect the the public repository decreases the security of PAS and your flows, both creating a security vulnerability and reducing the control over the changing of flows used in your production environment.

Configuring heap size for PAS server process

You can optimize performance of Central by configuring the maximum amount of random-access memory (RAM) that is allowed to the Central process. This value is not pre-allocated when the Jetty service starts.

To configure heap size

- In the PAS home directory, open the jetty subdirectory, and then open Start_jetty.bat for editing.
- In the line "SET ICONCLUDE_MEM_OPTS=-Xmx768m", change "768" to a value that represents a desirable maximum amount of memory for the service.
Be sure to append "m" to the value to specify that the value represents an amount in megabytes. (If you do not append the "m", the maximum memory is specified as bytes.)
- In the PAS home directory, navigate to jetty\win32 and open Wrapper.conf.
- Change the line "wrapper.java.maxmemory=768" to a value that represents a desirable maximum amount of memory for the service.

In Wrapper.conf, you can leave 'm' off the end of the value, because this value always represents megabytes.

These two values (in Start_jetty.bat and in Wrapper.conf), should be the same, but they do not have to be. Start_jetty.bat starts Central as a stand-alone, command-line process, and wrapper.conf controls Central when started as a Win32 service.

Configuring log file settings

PAS records errors (ERROR), warnings (WARN), information (INFO), and debugging messages (DEBUG) in the following log files:

- For Studio: iConclude.log, in the \Studio\logs subdirectory of the PAS home directory
- For Central: Wrapper.log, in the \Central\logs subdirectory of the PAS home directory

Because logging activity can slow the PAS's performance and create very large log files, it is important that PAS run with appropriate logging levels. The default logging levels have been set to provide necessary information without impacting performance. It is recommended that you use the default logging levels.

To change logging levels

1. In the jetty\resources subdirectory of the PAS home directory, modify the log4j.properties file according to your needs.
2. Save changes.

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 PAS 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.