




SRM Suite

Version 3.0

User and Administration Guide

302-000-514

01

Copyright © 2013 EMC Corporation. All rights reserved. Published in USA.

Published December, 2013

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>). For documentation on EMC Data Domain products, go to the EMC Data Domain Support Portal (<https://my.datadomain.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

Tables		9
Chapter 1	Browsing Reports	11
	Browse mode	12
	Dragging a report to a new position	12
	Resizing a report	12
	Applying filters to table columns	12
	Toggling values in graphs	13
	Setting upper and lower bounds to graphs	13
	Reverting to the original bounds of a graph	13
	Changing the appearance of graphs	13
	Arranging reports in the report pane	13
	Changing the report type	14
	Displaying detailed graph data in a popup	14
	Zooming in on graph data	14
	Zooming in on data while displaying the complete time range	14
	Changing the sampling period and duration	15
	Displaying a report in a new browser window	15
	Creating a dashboard of frequently viewed reports	15
	Comparing reports in a Personal View	15
	Editing a report context in a Personal View	16
	Saving a favorite report	16
	Displaying a favorite report at login	16
	Storing a snapshot of a report	16
	Getting the complete URL for a report	17
	Printing a report	17
	Emailing a report to yourself	17
	Creating a scheduled report	17
	Launching a scheduled report now	17
	Searching for a report	18
	Exporting a report	18
	Creating a report with the Report Wizard	18
Chapter 2	Working with Reports in Edit mode	19
	Report workflow	20
	Filtering and expansion	20
	Defining the report metrics using a filter	21
	Syntax for filter expressions	21
	Specifying the expansion properties for sub nodes	22
	Complex node expansions	22
	Configuring the advanced parameters for a report node	26
	Setting the child node ordering	26
	Filtering and expansion parameters	26
	Report configuration	29
	Defining report types	29
	Report configuration parameters	29
	Report Details	31
	Table reports	32

	Graph reports.....	39
	TopN reports.....	43
	Map reports.....	46
	Heat map reports.....	48
	Treemap reports.....	49
	Editing external reports.....	51
	Topology reports.....	51
	Overlay reports.....	53
	Status reports.....	55
	Formula configurations.....	56
	Formula parameters.....	56
	Creating pre-generated reports.....	58
	Report URL syntax.....	58
Chapter 3	Managing ReportPacks	63
	Creating a ReportPack.....	64
	Uploading a <ReportPack>.arp file.....	64
	Exporting a ReportPack.....	64
	Deleting a ReportPack.....	65
	Adding a new report template to a ReportPack.....	65
	Copying a report template.....	65
	Deleting a report template from a ReportPack.....	66
	Generating a report template in XML format.....	66
Chapter 4	Managing Users	67
	How the interface manages user rights.....	68
	Adding a user.....	68
	Setting the access rights of a user.....	69
	Setting the viewing rights of a user.....	70
	Testing user settings.....	70
	Copying users.....	70
	Editing users.....	71
	Deleting users.....	71
	Viewing users.....	71
	Creating a role.....	71
	Adding members to a role.....	72
	Adding external members to a role.....	72
	Establishing inheritance between roles.....	72
	Setting role access to report templates.....	73
	Setting component access and restrictions.....	73
	Assigning administration tasks to a role.....	74
	Editing a role.....	74
	Deleting a role.....	75
	Removing a user from a role.....	75
	Resetting default role settings.....	75
	Viewing roles.....	75
	Creating a profile.....	76
	Customizing a profile.....	76
	Adding members to a profile.....	76
	Adding external members to a profile.....	77
	Editing profiles.....	77
	Removing a member from a profile.....	77
	Viewing profiles.....	78

Chapter 5	Administering the Interface	79
	Setting the login properties.....	80
	Customizing the look of the interface.....	80
	Enabling keyboard shortcuts.....	80
	Modifying the default profile settings.....	81
	Editing the aliases given to time periods.....	81
	Disabling user logins.....	81
	Cancelling a user session.....	81
	Accessing the scheduled reporting options.....	82
	Editing the properties of scheduled reports.....	82
	Listing reports that are generating.....	83
	Deleting stored reports.....	83
	Accessing modules.....	83
	Deleting metrics.....	83
 Chapter 6	 Customizing User Settings	 85
	Modifying your user data.....	86
	Setting your reporting preferences.....	86
	Saving the report tree.....	86
	Uploading a report.....	87
 Chapter 7	 Licensing	 89
	Adding a SolutionPack license.....	90
	Adding licenses.....	90
	Verifying core software licenses from Centralized Management.....	91
	Obtaining a new license after host reconfiguration.....	91
 Chapter 8	 Ensuring Compliance	 93
	Policy Management.....	94
	Creating a policy.....	94
	Enabling or Disabling a policy.....	96
	Editing a policy.....	96
	Copying a policy.....	97
	Deleting a policy.....	97
	Notifications for compliance breaches.....	97
	Configuring alert definitions for notification of compliance breaches.....	98
	Setting a filter for compliance breach notification.....	98
	Enabling notifications for compliance breaches.....	99
	Disabling notifications for compliance breaches.....	99
	EMC Support Matrix Management.....	99
	Manual-matching to ESM.....	100
	Bulk-matching to ESM.....	100
 Chapter 9	 Alerting	 101
	Threshold alerts.....	102
	Prerequisites for generating threshold alerts.....	102
	Enabling threshold alerts.....	103
	Disabling threshold alerts.....	103
	Customizing an out-of-box threshold alert definition.....	103

	Creating a custom threshold alert definition for out-of-box SolutionPack support.....	104
	Availability alerts for PowerPath health metrics.....	104
	Prerequisites for generating availability alerts.....	105
	Disabling physical host availability alerts.....	105
	Enabling physical host availability alerts.....	105
	Disabling ESX host availability alerts.....	106
	Enabling ESX host availability alerts.....	106
	Alert notifications for consolidated alerts.....	106
	Prerequisites for generating alert notifications.....	107
	Tasks for enabling alert notification.....	107
	Disabling email notifications for consolidated alerts.....	109
	Disabling SNMP trap notifications for consolidated alerts.....	110
	Customizing an out-of-box alert definition for notification of consolidated alerts.....	110
	Creating a custom alert definition for notification of consolidated alerts.....	111
	Alert retention.....	112
Chapter 10	Online updates	113
	Configuring server settings for online updates.....	114
	Enabling the online update task.....	114
	Disabling the online update task.....	115
	Executing online update task manually.....	115
Chapter 11	Troubleshooting	117
	SSL certificates.....	118
	SSL certificate error examples.....	118
	Troubleshooting an SSL certificate error: Accepting the exception..	119
	Installing a self-signed SSL certificate from an Internet Explorer browser.....	119
	Importing an SSL certificate into your truststore.....	119
	Exporting an SSL certificate from an Internet Explorer browser.....	120
	Switch discovery failure in the SolutionPack for Cisco MDS/Nexus.....	121
	Troubleshooting service start-up problems on Windows.....	122
	Troubleshooting service start-up problems on UNIX.....	122
	Troubleshooting collector errors with Collector-Manager log files.....	123
	Licenses have expired.....	123
	Installing licenses.....	123
	Adding licenses.....	124
	Collect diagnostics on UNIX.....	124
	Collect diagnostics on Windows.....	125
	Troubleshooting alert collection errors on SolutionPack for EMC Data Protection Advisor.....	125
	Troubleshooting report generation errors on SolutionPack for EMC Data Protection Advisor.....	126
Chapter 12	Reference	127
	Report URL syntax.....	128
	Uploading your company logo.....	131
	Importing users.....	132
	Configuring the user process limits for the Linux installation.....	132
	Logging in to the User Interface.....	132

Starting the core software services from the command line.....	133
Starting the core software services on Windows.....	133
Verifying if the core software services are running.....	134
Connecting to Centralized Management.....	135
Stopping core software services on a UNIX server.....	135
Stopping core software services on a Windows server.....	135

TABLES

1	Table reports.....	32
2	Graph reports.....	39
3	Top N reports.....	43
4	Out-of-box Alert Consolidation Options During SolutionPack Installation.....	102
5	Mandatory configurations for email and trap alert notifications.....	108
6	Mandatory configurations for email and trap alert notifications.....	110
7	Mandatory configurations for email and trap alert notifications.....	111

CHAPTER 1

Browsing Reports

This chapter includes the following topics:

◆ Browse mode.....	12
◆ Dragging a report to a new position.....	12
◆ Resizing a report.....	12
◆ Applying filters to table columns.....	12
◆ Toggling values in graphs.....	13
◆ Setting upper and lower bounds to graphs.....	13
◆ Reverting to the original bounds of a graph.....	13
◆ Changing the appearance of graphs.....	13
◆ Arranging reports in the report pane.....	13
◆ Changing the report type.....	14
◆ Displaying detailed graph data in a popup.....	14
◆ Zooming in on graph data.....	14
◆ Zooming in on data while displaying the complete time range.....	14
◆ Changing the sampling period and duration.....	15
◆ Displaying a report in a new browser window.....	15
◆ Creating a dashboard of frequently viewed reports.....	15
◆ Comparing reports in a Personal View.....	15
◆ Editing a report context in a Personal View.....	16
◆ Saving a favorite report.....	16
◆ Displaying a favorite report at login.....	16
◆ Storing a snapshot of a report.....	16
◆ Getting the complete URL for a report.....	17
◆ Printing a report.....	17
◆ Emailing a report to yourself.....	17
◆ Creating a scheduled report.....	17
◆ Launching a scheduled report now.....	17
◆ Searching for a report.....	18
◆ Exporting a report.....	18
◆ Creating a report with the Report Wizard.....	18

Browse mode

Browse mode is the default mode of the user interface. In Browse mode, you can generate and view reports on the fly by clicking the report nodes in the tree.

You can do the following in browse mode.

- ◆ Drill down in most reports to get more detailed reports
- ◆ Change how reports are displayed, including the type of table or graph that is used to display the data
- ◆ Customize the look and feel of the displayed metrics
- ◆ Change the type of aggregation and the type range of the report
- ◆ Create a dashboard of favorite reports
- ◆ Take a snapshot of a report
- ◆ Print, email, and export reports
- ◆ Schedule reports

Reports are generated according to the report configuration parameters, which you can only edit while in Edit mode.

Dragging a report to a new position

You can move a report to a different position in the report pane.

Procedure

1. Mouse over the upper right corner of a report.
A set of icons appear.
2. Click and hold the **Drag this element to another place** icon and move the report.

Resizing a report

You can adjust the width and height of a report to accommodate your display by using the report resize icons.

Procedure

1. Mouse over the upper right corner of a report .
A set of icons appear.
2. Click a resize icon.
3. To save the change, click **save**.

Applying filters to table columns

Filtering a table column gives you only the values you want to see.

Procedure

1. Click the funnel icon above a column.
2. Enter a value.

You can revert to the original report by clicking **cancel**.

Toggling values in graphs

When metrics appear in the legend, you can toggle their display on the graph.

Procedure

1. Click the metric in the legend that you do not want displayed.
2. Click **Apply** to display the remaining metrics.

Setting upper and lower bounds to graphs

You can change the displayed range of values for a graph.

Procedure

1. In the **Lower Bound** field, enter the lowest value to display.
2. In the **Upper Bound** field, enter the highest value to display.
3. Click **Apply**.

Reverting to the original bounds of a graph

If you made changes to the bounds of a graph, use this procedure to revert to the original bounds.

Procedure

1. In the **Lower Bound** field, delete the value.
2. In the **Upper Bound** field, delete the value.
3. Click **Apply**.

Changing the appearance of graphs

You can customize the curve, width, marker, and color of graph metrics.

Procedure

1. Display a graph.
2. Click the **Graph Customization** icon.
3. Change the curve, width, marker, and color for the metric.
4. Click **Ok**.

Arranging reports in the report pane

By default, multiple reports are stacked one on top of one another. You can display multiple reports side by side.

Procedure

1. Click **Display**.
2. In the **Report Type** area, click **Display**.
3. Select how many reports to display per line.

4. Click **Apply**.

Changing the report type

You can choose other compatible report types to display data. If there is no compatible report type for the report, you cannot change the report type.

Procedure

1. Click **Display**.
2. In the **Report Type** area, click **Report**.
A list of compatible reports appears.
3. Select the report type.
4. Click **Apply**.

Displaying detailed graph data in a popup

With dynamic graphs you can mouse over a graph metric to display details about a selected point.

Procedure

1. Click **Display**.
2. In the **Report Type** area, click **Rendering**.
3. Select **Dynamic Charts**
4. Click **Apply**.
5. Mouse over the metric to display the time and value of a selected point.

Zooming in on graph data

With dynamic charts you can view graph details by enlarging an area of the graph.

Procedure

1. Click **Display**.
2. In the **Report Type** area, click **Rendering**.
3. Select **dynamic charts**.
4. Click **Apply**.
5. Use the mouse to select an area of the graph.
The graph zooms in to display the selected area.

Zooming in on data while displaying the complete time range

With dynamic charts you can zoom in on a period of time while also displaying the complete time range of the report.

Procedure

1. Click **Display**.
2. In the **Report Type** area, click **Rendering**.

3. Select **dynamic charts**.
4. Click **Apply**.
5. Click the magnifying glass icon.
A smaller graph is displayed.
6. Use the mouse to select an area of the smaller graph.
The graph zooms in to display the selected area.

Changing the sampling period and duration

You can set reports to display a variety of time period aggregates over a selected time range.

Procedure

1. Click **Display**.
2. In the **Time Selection** area, select the time period, aggregation, and duration of the collected data.
3. To lock the settings, click the lock icon.
4. Click **Apply**.

Displaying a report in a new browser window

You can view a report by itself in a browser window.

Procedure

1. Click **Display**.
2. In the **Actions** area, click **View in Full Page**.

Creating a dashboard of frequently viewed reports

Clicking on the pin icon of any report adds that report to the **My View** node in the report tree. By clicking on the **My View** node, you can view all of the pinned reports in one place.

Procedure

1. Mouse over the upper right corner of a report.
A set of icons appear.
2. Click the pin icon.
3. If you did not yet create a personal view, click **New Personal View** to create one. Otherwise, select the view in which to add the report.

Comparing reports in a Personal View

You can create a new report, based on an existing pinned one, for comparison purposes.

Procedure

1. Select the pinned report in a **Personal View**.
2. Click **Compare to new report** from the gear icon.

3. Type a name for the new report.
4. Make the appropriate changes.
5. Click **Ok**.

Editing a report context in a Personal View

You can edit an existing pinned report for customization purposes. This includes the context, which are the filter properties applied to a report's parent and linked node.

Procedure

1. Select the pinned report in a **Personal View**.
2. Click **Report context** from the gear icon.
3. Make the appropriate changes.
4. Click **Ok**.

Saving a favorite report

Favorite reports are like bookmarks that give you quick access to reports you use frequently.

Procedure

1. Click **Tools > Favorite this Report**.
2. Click **Save**.

Displaying a favorite report at login

You can automatically view a favorite report in the report pane after you login.

Procedure

1. Click **Favorite Reports**.
2. Expand the report information of the report you want displayed at login.
3. Click **Use as login report**.

Storing a snapshot of a report

You can store a point-in-time copy of a report at a given time.

Procedure

1. Click **Tools > Store this report**.

A snapshot of the report is stored under **Stored Reports**. When one or more stored reports are available for a node, a drop-down is displayed after the report is generated which gives you the option of selecting a stored report. By default, stored reports are synchronized with nodes every hour or whenever you click on **Stored Reports**.

Getting the complete URL for a report

With the complete URL of a report you can email the link to a report or bookmark the report in your browser.

Procedure

1. Click **Tools** > **Show Report URL**.

The URL for the report appears in a popup.

2. Copy the URL.

Printing a report

You can print a hard copy of a report for your records.

Procedure

1. Click **Tools** > **Print this Report**.
2. Set the printer information and click **Print**.

Emailing a report to yourself

You can send a report as a PDF attachment to your email address.

Before you begin

Verify your email address in **Settings**.

Procedure

1. Click **Tools** > **E-mail me this Report**.

The report is sent to the email address specified in your user settings.

Creating a scheduled report

Scheduled reports are sent at fixed times to a list of recipients.

Procedure

1. Click **Tools** > **Schedule this Report**.
2. Enter a name for the report.
3. Set the time of the report.

Whenever a scheduled report is generated it queries the database for the latest data.

4. Click **Save**.

Launching a scheduled report now

You can launch a scheduled report at any time.

Procedure

1. In the report tree, click **Scheduled Reports**.
2. Click the report you want to run.

A popup appears.

3. Click **Launch now**.

Searching for a report

You can quickly access a report without browsing the report tree.

Procedure

1. In the **Search** box, type the keywords for the report, such as a device name or IP address.
2. Press **Enter**.
3. Click the link to display the report.

Exporting a report

You can export reports into PDF, CSV, XLS, PNG, JPEG, SVG, and XML formats.

Procedure

1. Click **Exports**.
2. Click an export format.

Creating a report with the Report Wizard

Table, graph, and TopN reports can be quickly generated using one of four predefined filters to select metrics: APG Topology, Core Networking, Device Browser, or Devices by Location.

Procedure

1. Click **Tools > Report Wizard**.
2. From **View Type**, select a view type.
3. In **Available Elements**, select the elements to filter on.
4. Click **Run**.
5. If you want to save the report to reuse it, click **Save in My Reports**.

CHAPTER 2

Working with Reports in Edit mode

This chapter includes the following topics:

◆ Report workflow	20
◆ Filtering and expansion	20
◆ Report configuration	29
◆ Report Details	31
◆ Formula configurations	56
◆ Creating pre-generated reports	58
◆ Report URL syntax	58

Report workflow

In Edit mode, you can create new reports, edit existing reports, and customize the report tree. Changes that you make in Edit mode are saved and available to all users.

There are two basic approaches to creating new reports. You can copy an existing report to use as the basis for the new report or you can create a new report from scratch. Either way, the workflow for creating reports is as follows:

1. Using the **Filtering and Expansion** tab, select the metrics to display in the report. Since the database may contain thousands of metrics, you must use filtering and expansion to narrow the metrics to display in a report.
2. Using the **Report Configuration** tab, set the type of report you want, such as a table or graph, and specify the time period and aggregation to use in a report.
3. Depending on the report type selected, the **Report Details: <Report Type>** tab appears where you set the parameters for presenting information in the report.
4. Using the **Formula** tab, you can apply formulas to any node of the tree to perform advanced computations on any number of values.

Working in My Reports

The **My Reports** area gives you a workspace for editing reports and testing changes before making the reports available to users. Always make a backup of the report tree before you begin making major changes to reports or the report tree. Before you edit a report, first copy it to the **My Reports** area. When you are done making changes, copy the report back to its location in the tree. A report copied to **My Reports** maintains no relation to the source report, and you can edit it without affecting the source report.

If you are using a fresh installation, the **My Reports** area does not appear until you create a report using the Create a new Report icon.

Filtering and expansion

Filtering and expansion affect the report node properties. These properties are crucial to every node and report type because they determine what data is retrieved from the database.

The name of the report node, the filter used for metric selection, and some advanced settings that you can adjust are on the **Filtering & Expansion** tab.

Filtering

The filter of a node is an expression used to select the metrics that make up the basic content of the report. Because the filter determines the report's data set, it is a vital part of the overall report configuration. To set the filter properly, you must know which properties are being collected by the collectors and what they mean. There are several ways to create a filter. If you know what properties and values you want to set, you can manually edit a filter. You can also use the **Filter Wizard** to help you define a filter.

Expansion

The expansion setting is central to how the report tree works. It enables the quick generation of report trees, expanding its branches as it transforms individual nodes into multiple nodes. Expansion can automatically create and delete nodes as they appear or disappear in the database.

Defining the report metrics using a filter

A filter narrows the scope of data to display in a report. You can edit the filter expressions created by the **Filter Wizard** or build your own expressions.

Procedure

1. In Edit mode, click the **Filtering & Expansion** tab.
2. In **Filter**, click **Everything > Refine > using a wizard**.

The **Filter Wizard** helps you build a filter to narrow the scope of data to display. The wizard can help you build complex filters that use logical operators.

3. Type the name of a property.

For example, typing **device** causes the **Filter Wizard** to display a list of known devices based on the data stored in the database.

4. Select a property value.

Each property has a value. For example, the value of the device property is the device name.

5. In **Use the filter of this node for**, specify how to use the filter.
6. In **Display Policy**, select how you want to display nodes.
7. Click **Save**.

Syntax for filter expressions

You can build your own filter expressions using predicates and operators that are grouped together by parenthesis. Before creating your own filter expressions, practice with the Filter Wizard.

Predicates

Predicate	Description
property='pattern'	Selects a variable based on a pattern. This pattern is a string that can contain SQL wildcards, such as the percent sign (%) and underscore (_).
property='value'	Selects a variable with the exact value match.
property	Selects only the variable of the specified property.
#<database type>-<database name>:id	Selects the variable with the ID in the specified database. When using a cache group in the context of the APG-Property-Store, the database name is the cache group name. For example: #APG-DB:ALL For events, use the name of the corresponding resource in the APG.xml file with the FLOW-keyword. For example: RPE2:ALL
#<database type>-<database name>:ALL	Selects every variable in the specified database. For example: #APG-DB:ALL
#<database type>:ALL	Selects every variable in the specified database type. For example: #APG:ALL

Predicate	Description
*	Indicates no filter is applied and everything in the database is available for filtering.

Operators

Operation	Operator
AND	&
OR	(pipe)
NOT	!
Is	=
Strictly is	==

Guidelines

Follow these guidelines when building and editing filter expressions:

- ◆ Enclose property values in single quotes. For example: `devtype=='Router'`
- ◆ Use the AND or OR operators to search for two or more property values. For example: `devtype=='Router' & parttype=='interface'`
- ◆ Use parenthesis to enclose predicates and operators. For example: `device=='Router_A' & (parttype=='Interface' | parttype=='Port')`

Specifying the expansion properties for sub nodes

Expansion dynamically creates sub nodes under a report based on the properties you select. When a node expands, the node transforms itself from a single node into as many nodes as there are matches of the specified property name.

You set expansion in Edit mode and display the expanded nodes in Browse mode.

Procedure

1. In Edit mode, click the **Filtering & Expansion** tab.
2. Expand **Expansion**.
3. Click **Add a Property**.

The **Property Selection Helper** contains one or more tabs of properties.

4. Select one or multiple properties from the tabs.

The properties selected display in the **Selected Properties** field.

5. Click **OK**.

Complex node expansions

Complex expansions perform calculations within the context of an expansion.

With complex node expansions, you can:

- ◆ Separate properties, each having a list of values.
- ◆ Group and remove nodes based on matching criteria.

- ◆ Remove filters applied to parent nodes.

Complex expansion types

You create complex expansions using four expansion types:

- ◆ dual-regex
- ◆ sql
- ◆ regex
- ◆ split

The expansion types create groups based on a subpart of a property value.

Expansions use pattern matching and transformations. The syntax for complex

expansions is: `property<type=complex expansion`

`type;param1=value1;param2=value2;...>`, where the expansion type is a string to identify which complex expansion to use. You cannot use the semi-color (;) or greater than (>) characters in a parameter value.

Dual Regex Pattern Match (type=dual-regex)

This expansion uses these parameters:

value-match	The first regular expression matched against the property values.
value-replace	The replacement string applied to the previously matched values. The result of this replacement is the group name.
group-match	A regular expression matched against the previously generated group names.
group-replace-regex	A replacement string to create a regular expression from the group name, to match property values include in this group.

The expansion occurs in these phases:

- ◆ For each distinct property value, matches the property value with `value-match`. If the value matches, replaces it using the replacement string `value-replace`. The resulting value creates a new node in the tree.
- ◆ For each node created, applies a transformation on its name using the regular expression `group-match` and replacement string `group-replace-regex`. This transformation creates a third regular expression.
- ◆ Rematches any property value with this regular expression. This works similar to a filter. If the value matches, it is selected by the group node.

For example, suppose there are three devices in a network expanding the property device into three values: dev-w4n-montreal, router-w4n-montreal, and server-w4n-toronto. To use the location name to group these devices, do the following:

- ◆ Create location groups from the device names with this expression: `^.*-w4n-(.*)`
`$ → $1`
- ◆ To match the groups montreal and toronto, use a matching device. A matching device is a device whose name ends with the group name, which is the city name in this example. For the montreal group, the matching regular expression is `^.*-montreal$`.
- ◆ To match a device in a group, build the previous regular expression from the group name with this replacement: `^(.*)$ → .*-\\Q$1\\E`

- ◆ The result is this expansion: `device<type=dual-regex;value-match=^.*-w4n-(.*)$;value-replace=$1;group-match=^(.*)$;group-replace-regex=.*-\\Q$1\\E>`

Combined Regex/Apg Pattern Match (type=sql)

This expansion uses two regular expressions and transformations to create new groups based on a pattern. This is similar to the Dual Regex Pattern Match except that the values are matched in groups using a filter instead of a regular expression. Since this expansion is quicker than the others, use this one whenever possible.

This expansion uses these parameters:

value-match	The first regular expression matched against the property values.
value-replace	The replacement string applied to the previously matched values. The result of this replacement is the group name
group-match	A regular expression matched against the previously generated group names.
group-replace-sql	A replacement string to create a pattern (with % and _ wildcards) to match property values to include in this group.

The expansion occurs in these phases:

- ◆ For each distinct property value, tries to match the property value with value-match. If the value matches, then replaces it using the replacement string value-replace. The resulting value creates a new node in the tree.
- ◆ For each node created in the first phase, applies a transformation on its name using the regular expression group-match and replacement string group-replace-sql. This transformation creates a valid pattern that is applied to the created node as `property='pattern'`.

For example, suppose there are three devices in a network so that the property device is expanded in three values: dev-w4n-montreal, router-w4n-montreal, server-w4n-toronto. To use the location to group devices, do the following:

- ◆ Create location groups from the device names with this expression: `^.*-w4n-(.*)$ → $1`
- ◆ To match the groups montreal and toronto, use a matching device. A matching device is a device whose name ends with the group name, which is the city name in this example. For the montreal group, the matching filter is `'device=%-montreal'`.
- ◆ The expansion result is: `device<type=sql;value-match=^.*-w4n-(.*)$;value-replace=$1;group-match=^(.*)$;group-replace-sql=%-$1>`

Regex Pattern Match (type=regex)

This expansion type is a simplified version of the two previous types. Most of the time, the values used to create group names belong to those groups, and only to those groups. This is what the Regex Pattern Match does, skipping the post matching phases of the Dual Regex Pattern Match and the Combined Regex/Pattern Match.

This expansion uses these parameters:

value-match	The first regular expression matched against the property values.
value-replace	The replacement string applied to the previously matched values. The result of this replacement is the group name.

For example, suppose there are three devices in a network that expands the property device into three values: dev-w4n-montreal, router-w4n-montreal, and server-w4n-toronto. To use the location to group devices, do the following:

The expansion occurs in these phases:

- ◆ For each distinct property value, tries to match the property value with value-match. If the value matches, then replaces it using the replacement string value-replace.
- ◆ The resulting value creates a new node in the tree. The value transformed into the group name is automatically included in the group.
- ◆ The expansion result is : device<type=regex;value-match=^.*-w4n-(.*)\$;value-replace=\$1>

Split Property Expansion (type=split)

This expansion type creates reports using a relation between objects. A relation stores one or several object names, such as a device name, in another object's property, such as the remote system property.

This expansion uses these parameters:

value-separator	The character used to separate entries in the targeted property.
property-separator	The character used to separate properties when values contain multiple entries
properties	The name of the properties whose values are extracted by the split operation. You can use multiple occurrences if more than one property needs to be identified in the generated filter.
level-up	Parent nodes to be skipped in the construction of the node's filter.
name-override	You can use a specific name for the created node.

In this example, the collected data includes the property peerifs that identifies peer interfaces. It contains a list of device-part pairs expressed like this:

```
peerifs='device1,part1;device1,part2;device2,part1'
```

To create a node for each distinct device-part pair and associate a filter to each node that selects the correct data: device=='device1' & part=='part1'. The result is:

```
peerifs<type=split;value-separator=~;;property-separator=,;properties=device;properties=part;level-up=5>
```

You can use the tilde (~) character to escape the semi-column when it is the targeted character to split values or properties. You can use the properties parameter several times if you want the filter to contain multiple parts. The name of the resulting node is the name of the matched value if not overridden by the name-override parameter.

In this example, you can use the split expansion to void the filtering effect of parent nodes. This is useful when the filter is incompatible with the data you want to report on. For example, to report on interface utilization (parttype='Interface') of devices with a Filesystem (parttype='Filesystem'), split on a property, create the filter using the same property, and use the level-up parameter to void the effect of the selection filter: device<type=split;properties=device;level-up=1>.

Configuring the advanced parameters for a report node

Advanced parameters provide features that enable you to create complex custom reports.

Procedure

1. In Edit mode, click the **Filtering & Expansion** tab.
2. Expand **Advanced Parameters**.
3. In **Use filter on unmatched variables**, select how variables that are not selected will be treated.
4. In **Report Display Preference**, select how the report is generated.
5. Select the node icon for the report node.
6. In **Top N for events**, enter the limit of results to return.
7. Click **Save**.

Setting the child node ordering

You can order child nodes either by position or alphabetically.

Procedure

1. In Edit mode, click the **Filtering & Expansion** tab.
2. Expand **Child Node Ordering**.
3. Drag and drop the children to either the **Ordered by Position** or **Ordered by Name** lists.
4. Click **Save**.

Filtering and expansion parameters

You define the scope of data to include in a report on the **Filtering & Expansion** tab. You define the basic properties of a report, such as its name, ID, icon, and when and how to display the report in the tree.

Name

The name of the report that you are creating. If you use expansion, the report name represents the expansion evaluation.

Unique identifier

You can link between reports by using the identifier. When you paste a report as a link, you can use the unique identifier if it exists. A unique identifier consists of letters, numbers, and the characters period (.) comma (,) dash (-), underscore (_) or spaces. You can also use report identifiers to define report restrictions and to control branches to be skipped by the frontend search engine.

This option is hidden by default. To access this option, click the anchor icon to the right of the **Name** field.

Filter

A filter is an expression that limits the number of metrics displayed on a report. A filter is vital to configuring a report because it determines the report's data set. You can manually create a filter when you know which properties and values to include in a filter, or you can use the **Filter Wizard** to help you define the filter.

Use the filter of this node for

Use this field when using both expansion and a filter.

expansion only	Applies the filter while expanding the node using expansion. Any nodes that do not match the filter are not created. Selected variables are not affected by the filter.
selection only	Applies the filter after node expansion. The system performs expansion without taking the filter into account, and passes each resulting node through the filter.
expansion and selection	Combines both expansion and selection. The filter of the node is used for node generation and each resulting node for the variables it selects.

Display Policy

Use this field to simplify the presentation of the report tree by hiding certain report nodes and showing only relevant report nodes.

Always show the node	Always shows the node in the report tree. This is the default behavior.
Hide if no variable is selected	Hides the node if its filter does not select any metric.
Hide if the node is a leaf	Hides the node if it does not have any children. This is particularly useful with expanded nodes, as they may expand into an empty list.
Hide if no variable is selected or if the node is a leaf	This is a combination of the Hide if no variable is selected and Hide if the node is a leaf policies.
Hide in Browse mode but use for report computation	This policy is useful when you define a node to build reports above it in the tree, but not as a report itself. When you chain formulas, use this policy.
Always hide this node	Never shows the node in the report tree. This can temporarily disable a report and any formulas applied to this node.
Hide except when it's the target of a link	Hides the node except when it is the target of a link.
Show this node only in the tree and not in the report	Forbids drilling down to this node from a parent report.

Expand on every

This setting is central to how the report tree works. Expansion dynamically creates sub nodes. Sub nodes add levels to a report so you can drill down. When a node is set to expand, the node transforms itself from a single node to as many nodes as there are matches of the specified property name in the data set.

By using expansion, you can define a report that automatically creates and deletes nodes based on the properties found in the data set. For example, as devices are added to the network, they are automatically added to the report when their properties match the expansion criteria.

Use filter on unmatched variables

Use this field to create nodes using variables that are not selected by the sibling nodes.

No	The node acts as an Others or Unmatched node.
-----------	---

For the edit node filter	Selects all the variables that are not selected by the sibling nodes with reference to the Edit mode filters. For example, if a few sibling nodes select the same IP address, the Other node selects all the other IP addresses. If there is no expansion, there is no distinction between the Edit and Browse node filters. Use this setting in most cases. If there is an expansion, this setting returns the unmatched results of the sibling filters as they appear in Edit mode. This differs from how the resulting node filters are expressed in Browse mode, and provides different results.
For the resulting browse node filter	Use this setting when using the Top N for the events property to select a reduced set of results in the expansion, or when doing complex expansions that result in a partial list of results. In both cases, the unmatched results are evaluated against the filters of the expanded nodes in Browse mode, resulting in all variables not chosen by any expanded nodes. In both cases, the results are subject to any parent node filters. You can use filters on these other variables that the node with the unmatched option selects.

Report Display Preference

You can use a stored report instead of creating a report that takes a considerable amount of time to generate.

Use Real Time Values	The report is generated on-the-fly.
Prefer Stored Report	When a user clicks on the report node, the most recent and available stored report appears. Otherwise, the report is generated on-the-fly.
Use Only Stored Report	When a user clicks on a report node, the most recent and available stored report appears. Otherwise, a notice appears indicating the system cannot display the report.

Node Icon

You can use a node icon for the following report areas:

- ◆ As a small icon in the tree alongside the node name
- ◆ As a medium icon on the report alongside the element titles
- ◆ As a big icon in icon reports

If you want an icon to be chosen automatically, use the default icon.

Top N for events

You can limit the number of results that are returned, which can significantly speed up the time it takes to generate reports. This setting is only applicable for events, and is only visible if it was declared in at least one of the event mapping files being used.

Child Node Ordering

You can arrange child nodes manually by dragging and dropping them in **Ordered by Position**, or you can arrange them alphabetically in **Ordered by Name**. You only can see this field if child nodes exist.

Report configuration

Use the **Report Configuration** tab to define the core settings for a report, including the report type, the time period of the report, and the aggregation used.

The settings on this tab deal with the report preferences that are used by default when the corresponding node is selected in Browse mode.

Defining report types

A report type determines the layout of a report. You select the report type on the **Report Configuration** tab.

Procedure

1. In **Display Mode**, indicate how you want multiple reports to appear on the report page.
The default behavior displays one report after another report.
2. In **Report Type**, click on the layout to use for the report.
3. In **Report Rendering**, select the type of rendering to apply to the report.
4. In **Default Duration**, specify the time range to display on the report.
5. In **Sampling Period**, indicate the time interval to use to collect data from the database.
6. In **Sampling Type**, specify the aggregation to use.
7. Click **Propagate time selection on drill-down** to apply the report's time periods to its sub-reports.
8. In **Report Description**, type the information to display at the top of the report. You can use XHTML in <div> tags in the description.
9. In **Displayed Properties**, click **Add property** for each property that you want to include in the report header. You can select one from a list of default properties or type a custom property.

Report configuration parameters

You set the type, the time period, the rendering, the aggregation, and the description of a report on the **Report Configuration** tab. You can also use this tab's advanced settings to define the time zone and indicate whether to use dynamic maintenance periods.

Display Mode

Indicate how to place multiple reports on the report page for viewing. The default behavior displays one report after another report. You can also manually rearrange reports on a page.

Report Type

Select which report type to use for displaying metrics on this node.

Report Rendering

Select the rendering to apply on the report, including any displayed sub-report.

Dynamic rendering is available on standard charts and provides better zooming and analyzing tools than the static one. However, it can have a performance impact if enabled on many reports simultaneously.

Default Duration

Select the time range to display on the report. Sometimes, this time range is not taken into account. For example, if the report is using a formula where the duration is explicitly

selected, this default value is bypassed. You can exclude time periods from the report by clicking the calendar icon.

Sampling Period

The aggregate interval to use to gather data from the data base. You can also configure the report to use aligned data for one of the aggregates.

Sampling Type

The aggregation to use to gather data from the database. This does not apply to some report types, such as pie charts and horizontal bars.

Propagate time selection on drill-down

Enable the temporal settings of the current report to use for all child reports when drilling down to them from the report in the page area. The settings are:

- ◆ Default Duration
- ◆ Sampling Period
- ◆ Sampling Time

Report Description

This description displays at the top of a report. You can use XHTML in <div> tags in the description.

Displayed Properties

These properties display at the top of the report, which include the contact name and location. The default name of the property can appear in its title, which displays in the gray text box.

Always use this setting when the report is for a single device because only one value appears for each added property.

Time Zone

Specify the time zone for a report generated on the current node. **Inherit** means that the time zone is inherited from the parent node. If there is no parent, the report uses the time zone from the profile of the current user.

The first field indicates whether the report uses universal time, the default, or whether the time period is aligned. Universal time indicates that the time period corresponds to the report's time zone in the adjacent list, and the data shown in the report is relative to its time zone. For example, selecting the last hour of data displays data from the current time, which is 10:00 America/Montreal time, and appears along the x-axis, back to 9:00 America/Montreal time. If you are monitoring a device in Montreal, you can see its activity data between 9:00 to 10:00. If you are also monitoring devices in Paris at this time, the activity data is between 3:00 and 4:00 AM Europe/Paris.

The other effect of universal time is that daily and weekly aggregates are delineated on UTC time. For example, the daily aggregation period for America/Montreal is from 18:00 - 19:00 whereas the Montreal time zone offset is UTC - 5 hours.

Align data on UTC using tz property with the following reference

This setting time shifts the data from different time zones to align them. The setting aligns data in order for the time zone to be independent in the report. Aligned data displays using aligned aggregates.

The effects of aligned time zone aggregates are:

- ◆ Data from different time zones is time-shifted to a universal day where data is expressed according to the local time zone where the data was recorded. For example, data at 9:00 locally is aligned with data at 9:00 at a remote time zone. Data can then be compared at each time zone in relation to how users experience it during a typical day or week. Without using aligned aggregates, data from different time zones is expressed in relation to the local time zone of the report.

- ◆ For aligned daily and weekly aggregates, days and weeks are delineated according to 24 hour 7 day periods where time zones are aligned. Without using aligned aggregates, days and weeks are delineated according to UTC.

When you use aligned data, the second time zone list sets the reference. This is important when choosing data from the last day. If the data is aligned and America/Montreal is the reference point, and it is 13:00 in Montreal and you display the last day of information, the data returned is from 13:00 today to 13:00 yesterday. If you choose Europe/Paris as the reference, and it is 19:00 in Paris, the last day extends the report from 19:00 on the current day to 19:00 yesterday and there won't be information in the report for the last 6 hours in Montreal because it has not occurred yet in the day where the time zones are aligned.

Reference is important if time zones occur on different days. The day the report covers is determined by the reference time zone.

For this functionality to be available, it must be configured in the backend and the frontend. Whether there are daily and weekly aligned aggregates is determined in the backend setup. For aligned aggregates of one hour or less, the frontend computes them on the fly on a best effort basis if aligned time zones are enabled. Some aggregates are not generated: the local time of the aggregate must be a number of whole hours offset from UTC for the one hour aligned aggregate to be computed. Locations with an additional half-hour offset such as Asia/Kolkata (UTC +5:30), are not computed.

You are not explicitly informed whether time zone aligned aggregates are available, but in most cases, if Align data on UTC using tz property with the following reference is available in the first time zone list, aligned time zone aggregates are available for each sampling period.

These constraints apply to aligned data:

- ◆ Maintenance periods configured in the interface appear in graphs but not used for computations or in tables. If you configure a maintenance period from 8 to 9, this period is highlighted in the graph, and pertains to 8 to 9 in all time zones, as the data is aligned.
- ◆ Do not use the Outage Editor with aligned data as the results are inaccurate.
- ◆ Do not plot events on the same graphs as aligned data with an Overlay, as the results are inaccurate. Events cannot be aligned across time zones as time series data can.

Dynamic Maintenance Periods

Set the maintenance period to use for this node according to the one set in the Outage Manager. Maintenance periods are incompatible with aligned time zones.

Report Details

Use the Report Details tabs to define the settings for a specific report type.

The Report Details tab changes according to the type of report you select on the **Report Configuration** tab. The following sets of report details are available:

- ◆ **Report Details: Table**
- ◆ **Report Details: Graph**
- ◆ **Report Details: TopN Graph**
- ◆ **Report Details: Map**
- ◆ **Report Details: Heat Map**
- ◆ **Report Details: Treemap**

- ◆ **Report Details: External**
- ◆ **Report Details: Topology**
- ◆ **Report Details: Overlays**
- ◆ **Report Details: Status**

Table reports

Table columns display information and table rows are populated by sub nodes. You can add columns to the table and change the configuration of the sub nodes that populate the rows.

Table columns display node attributes, properties, and values. You can add new columns and change the order of the columns.

Table reports include one row for each sub node. Usually, the sub node is expanded on a property.

For a report category all of the report details are the same. This means that you can change the report type to another one of the same category and the report settings remain applicable. After you save the settings, the report details are saved from session to session, even if you switch to another report type.

In **Report Type** on the **Report Configuration** tab, you can select these table reports.

Table report types

In **Report Type** on the **Report Configuration** tab, you can select these table reports.

Table 1 Table reports

Type	Description
list	Lists the child node names of a parent. You can drill down to more detailed reports. This is the default report type.
standard table	Displays one row per child node. You select the columns, such as availability. You then expand the table node with child nodes in the report tree, which automatically generates the data in the rows of the report. For example, if a child node expands on host, each row contains data about individual hosts in the network.
metric-based table	Displays the metrics of the child nodes.
Interactive table	Controls the display of adjacent reports. You can navigate from one report to another while maintaining a list of reports on the same page. The reports that display next to the interactive table are the child nodes of the interactive table. You can place these tables in mixed reports.

Setting major and critical thresholds

Set thresholds to alert you when a metric has exceeded a major or critical threshold.

Procedure

1. Click the **Report Details: Table** tab.
2. Expand a value.
3. Expand **Value Settings**.
4. In **Critical Level**, enter the value of the critical threshold.

5. In **Major Level**, enter the value of the major threshold.
6. In **Value Formatter**, select how exceeded thresholds will be displayed in the report.
7. Click **Save**.

Setting the number of rows to display per page

For a table with multiple pages, you can set the number of rows to display per page in the report pane.

Procedure

1. Click the **Report Configuration** tab.
2. Expand **Advanced Settings**.
3. In **Paging Limit**, enter the number of rows to display per page.
4. Click **Save**.

Adding a column

You can add details to tables by adding columns for the kind of information you want to add, such as node attributes, metric properties, and metric values.

Procedure

1. Click the **Report Details: Table** tab.
2. Click the **Attribute**, **Property**, or **Value** icon to add a column.

Option	Description
Attribute	Information about the node such as the node ID, node name, node filter and so on.
Property	Metadata about a metric, such as device, device type, group and so on.
Value	Raw value of a metric.

3. In **Column Name**, enter a name for the column.
4. Select the attribute, property, or value to display in the column.
5. Click **Save**.

Moving a column

You can change the order of the columns.

Procedure

1. Click the **Report Details: Table** tab.
2. Drag and drop a column header to a new location.

Copying a column

You can use column definitions that are already defined as part of a new table report.

Procedure

1. Click the **Report Details: Table** tab.
2. Check the columns that you want to copy.
3. Click **Copy**.
4. Navigate to another table report.

- 5. Click the **Report Details: Table** tab.
- 6. Click **Paste**.
- 7. Click **Save**.

Deleting a column

You can delete a table column.

Procedure

- 1. Click the **Report Details: Table** tab.
- 2. Click on the trash icon in the column header.

Table report parameters

You define the metrics for the columns in a table report on the **Report Details: Table** tab.

Attribute

You define the lines of a standard table report using the direct children of the selected node. An attribute column displays the information related to the current parent node, such as its name and number of children.

Column Name	A unique name that appears in the table header.
Attribute	◆ Other displays the name for the property you type in the field box. ◆ Node ID displays the node short ID relative to its parent. ◆ Global ID displays the complete node ID starting from the root of the tree. ◆ Node name displays the name of the node. ◆ Global name displays the complete name of the node, starting from the root of the tree. ◆ Expansion names array is only applicable if a multi-expansion is applied to the node that appears in the table. Type the index of this zero-based array that contains each part of the multi-expansion. ◆ Node filter displays the node short filter, relative to its parent. ◆ Global filter displays the complete node filter, starting from the root of the tree. ◆ Child count displays the number of children of that node.
Display Condition	◆ Show displays the column. ◆ Hide hides the column. ◆ If you specify a condition, the column appears when the current node's name matches the regular expression.

Property

A property column displays the property about the selected variables of a child node. This can be any collected property, such as a server name for a device.

Column Name	A unique name that appears in the table header.
--------------------	---

Property	Type the name of a property or click the Property Selection Helper icon and select one or more properties.
Selected Value	The system can retrieve a value from different locations. ♦ filter allows the configuration of a specific filter through the Filter Wizard, which selects the property to display from the child. ♦ formula result displays the property retrieved from the result of a formula applied on the children.
Filter to Apply	Refines the properties to display in this column. This filter applies to each line of the table, which represents each child node.
Use time settings from	Shares the time management parameters between columns whose result values come from the same formula. This setting may decrease the time to generate a report because the formula is evaluated only once.
Sampling Period	Selects the sampling that is set globally on the report, or enforces one from the computed aggregates available in the database. Adjusting this option can significantly reduce report generation time. For example, if you are reporting on a month or year, selecting a day or week aggregate instead of real-time dramatically increases the report generation speed. This is important because with such long ranges, a fine level of detail is generally not necessary. The period selection option is related to the sampling period. This is because the application can automatically select a higher aggregate than the one selected when it really fits the report time-range. This provides better performance and scales no matter the report time range. When you are unable to change the specified aggregate, always choose this option. However, this can trigger a high report generation time if the report time range is too wide. Note When you select a small sampling period, such as real-time, allow for the selection of higher periods for better performance.
Sampling Type	Selects one of the available aggregated values stored in the database. You can choose average, min, max, sum, the last value, count, the number of received values, or last timestamp, and the timestamp of the last received value, on the previously selected aggregate period.
Column Time Range(s)	Selects whether to use the global report time range for the values of this column, or a fixed one. A fixed range is relative to the end time of the report. It may also be divided into time slices, creating several columns in Browse mode.
Recover...	Determines if all values are retrieved over the report time range, or only the last value. If all values are retrieved, then they are aggregated to a single value using the selected function in the drop-down list.
Time Threshold	The time interval relative to the report time value, which acts like a tolerance interval for data retrieval. If you select real-time, the threshold is in seconds. If you select an aggregate sampling type, the time threshold is a number of periods.

	Note If you set zero for a time threshold while recovering only the last value using a real-time sampling period, an empty table cell may result. Always set a meaningful threshold according to the sampling period.
Temporal Aggregation	Aggregates the results of a formula based on the function that you select.
Display Condition	◆ To display the column, select show, the default. ◆ To hide the column, select hide. ◆ If you specify a condition, the column appears when the current node's name matches the regular expression.

Value

A value column retrieves a value from the database. You specify the value with a filter, such as CPU utilization, and the format for the value, such as percentage or bits per second.

Column Name	A unique name that appears in the table header.
Selected Value	Retrieves a value from the following: ◆ filter allows the configuration of a specific filter through the Filter Wizard, which selects the value to display from the child. ◆ formula result displays the value retrieved from the result of a formula applied on the children.
Filter to Apply	Refines the values to display in this column. This filter applies to each line of the table, which represents each child node.
Sampling Period	Selects the sampling that is set globally on the report, or enforces one from the computed aggregates available in the database. Adjusting this option can significantly reduce report generation time. For example, if you are reporting on a month or year, selecting a day or week aggregate instead of real-time dramatically increases the report generation speed. This is important because with such long ranges, such detail is not typically required. The period selection option is related to the sampling period. This is because the application can automatically select a higher aggregate than the one selected when it really fits the report time-range. This provides better performance and scales no matter the report time range. When you unable to change the specified aggregate, always choose this option. However, this can trigger a high report generation time if the report time range is too wide. Note When you select a small sampling period, such as real-time, allow for the selection of higher periods for better performance.
Sampling Type	Selects one of the available aggregated values stored in the database. You can choose average, min, max, sum, the last value,

	count, the number of received values, or last timestamp, and the timestamp of the last received value, on the previously selected aggregate period.
Column Time Range(s)	Selects whether to use the global report time range for the values of this column, or a fixed one. A fixed range is relative to the end time of the report. It may also be divided into time slices, creating several columns in Browse mode.
Recover...	Determines if all values are retrieved over the report time range, or only the last value. If all values are retrieved, then they are aggregated into a single value using the selected function.
Time Threshold	The time interval relative to the report time value, which acts like a tolerance interval for data retrieval. If you select real-time, the threshold is in seconds. If you select an aggregate sampling type, the time threshold is a number of periods. Note If you set zero for a time threshold while recovering only the last value using a real-time sampling period, an empty table cell may result. Always set a meaningful threshold according to the sampling period.
Use Rounding	For float values, you can specify how many figures to have after the decimal point by checking this box. In the Rounding Number box, type the number of decimal places to show in the value. For example, if you enter 2 in the field, a value of 3.14159 is displayed as 3.14.
Scaling Mode	If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.
Critical Level	The value for the critical threshold level. You can leave this empty to deactivate this level.
Major Level	The value for the major threshold level. You can leave this empty to deactivate this level.
Value Display	You can hide lines whose values are below a threshold by making a selection from this list. You can choose to display all values, only major and critical values, or just critical values. This enables you to generate exception reports that present only relevant data.
Is criticality ascending?	If one of the levels is empty or both levels have the same value, select or clear that setting according to your requirements. This option indicates whether the criticality is ascending where a higher number is more critical, or descending where lower values are more critical. If a different number is given for both levels, the direction of criticality is obvious and you do not need select or clear this option.
Value Formatter	Controls the display value for the column. ◆ Status displays the status icon for the value according to the threshold levels that are set. ◆ Value displays the value in numeric form and the status icon. ◆ Value with background displays the value with the background color matching the defined thresholds.

	◆ Value with graph displays the value and status icon, and is clickable so you can display an inline graph pop-up that plots the value over the time period. ◆ Date and time displays the date and time when the value was collected, converted from the UNIX timestamp, along with the status icon. ◆ Time displays the time when the value was collected, converted from the UNIX timestamp, along with the status icon. ◆ Duration formats a number of seconds into a duration. ◆ Sparkline displays the value and status icon and an inline sparkline, plotting the value over the selected time period. The dot is the position of the value on the graph, to give an idea of the scale, and the color corresponds to the threshold. ◆ Status sparkline displays the value and status icon and an inline status bar showing the status value over the time period. The line shows the status of each point of a value. If the band is green, the value is OK, orange indicates a problem, and red is critical. ◆ Performance sparkline displays the value and status icon and an inline performance bar showing the performance of the value relative to the other values on the page for this column, over the time period. The performance sparkline is like a gauge. It shows the value and where it is located compared to its thresholds. ◆ Bar chart sparkline displays the value and an inline bar graph whose number of bars depends on the selected sampling period and report time range. For this formatter, you can specify the bars' width and spacing in pixels. You can use a time zone from the report, inherited from a parent, or set from the user's profile.
Display Condition	Selects the display of the column. If specifying a condition, the column displays if the current node's name matches the regular expression.
Complement value	Converts a value to its complement, such as $1 - \text{value}$. Values are converted to their complements only after passing through all other processing, such as scaling.
Replace empty values by zero?	Shows zero for values that are empty. Values are replaced by zero only after passing through all other processing.
User per-line maintenance period?	Uses the maintenance period defined on each child node. The default uses the maintenance period for the current node, but that may not reflect the real maintenance period for every child.
Allow filtering on the value?	Triggers the use of a JEXL Boolean expression in which the value must match in order to display. The $\\$v$ variable represents the numeric value. In the following examples, rows are only shown when: ◆ <code>empty(v)</code> : the value is empty ◆ <code>(v >= 0.00 or v < 0.00) and not empty(v)</code> : the value is (greater or equal to 0.00 or less than 0.00) and not empty

- ◆ `v > 10.00 or v < 5.00`: the value is greater than 10.00 or less than 5.00
- ◆ `v != 0.00 and not empty (v)`: the value is not equal to 0.00 and not empty
- ◆ `v == 0.00 && not empty (v)`: the value is equal to 0.00 and not empty

Paging mode

Indicates the number of lines to display on one page of a table.

- ◆ **Show all values:** all pages are available.
- ◆ **Do not show if more than N values:** nothing displays when there are more values than what is entered for the Paging limit.
- ◆ **Show the first N values:** displays only the number of values for the Paging limit.

Paging limit

Defines the length of a table's page according to the values for the Paging mode.

Graph reports

You can add more graphing information or precision to any graph report, including standard, aggregated, and baseline graph reports.

Graph reports plot variables along the X (time) and Y (value) axes. Depending on the graph type, metrics can come from the current node or from the children nodes. You can use time series graphs that display multiple metrics related to a data selection or events that show only one series of values per event series.

There are value and baseline graphs that are available on every node of the tree, although baseline graphs require more than one week of data. Value graphs are generated using the properties of a selected node and its filter.

Graph reports can also have two y-axes, with two different sets of metrics plotted against them.

Graph report types

In **Report Type** on the **Report Configuration** tab, you can select these graph reports.

Table 2 Graph reports

Type	Description
simple chart	Gathers variables from the selected node. Draws a single graph showing variables as points over a time series.
bar chart	Gathers variables from the selected node. Draws a single graph showing variables as bars over a period of time. It is compatible with a simple chart and baseline reports.
children-based chart	Gathers variables from the child nodes of a selected node. Draws a single graph with one metric per child node. It is compatible with the stacked chart and the stacked bar.
per children charts	Gathers variables from each child node to create a simple chart, and then combines all charts into a single report.

Table 2 Graph reports (continued)

Type	Description
hierarchy charts	Gathers variables for each node in the tree to create a simple chart, and then combines all charts in a hierarchy that mirrors the tree structure.
stacked chart	Displays the aggregate value of each child node in an area connected by time series points, and each child node is stacked on top of one other. For example, it can stack the traffic of each router, showing the total traffic of a network by interface. This report is compatible with the children-based chart and the stacked bar chart.
stacked bars	Displays values in vertical bars instead of connected time series points. This report is compatible with the children-based chart and the stacked chart.
gauge chart	Displays a snapshot of the health of an indicator. It uses a single aggregate value of a child node. You set the thresholds to color code the range of values for this measurement.
baseline	Displays two lines, black and colored, and a gray zone. The colored line is the average value of the selected metric. The black line is the baseline and the surrounding grayed zone represents the deviation from the minimum and maximum value of the baseline.

Setting major and critical thresholds

Set thresholds to alert you when a metric has exceeded a major or critical threshold.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Thresholds**, enter the major and critical threshold values.
3. Click **Save**.

Setting the Y-Axis range of values

The Y-axis of a graph can be changed to display a specific range of values.

Procedure

1. Click the **Report Details: Graph** tab.
2. In **Main Y-Axis Boundaries**, enter the values you want to make the main Y-axis of the graph.
3. Click **Save**.

Scaling the Y-Axis

If the unit of the Y-axis values do not correspond to what you would like to display, you can change the unit by multiplying or dividing the values by a specified factor.

Procedure

1. Click the **Report Details: Graph** tab.

2. Use **Main Y-Axis Scaling Mode** to multiply or divide all of the values of the main Y-axis by a specified factor.
3. Click **Save**.

Setting the appearance of the graph metrics

You can customize the color and shape of the metrics that are displayed in the graph.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Metric Customization**, click **Add Customization**.
3. Select the metric you want to customize.
4. From the configuration options, choose whether the metric should be displayed against the main or the secondary axis, and how you want the line to be displayed.
5. Click **Save**.

Changing the legend

You can configure which properties are displayed in the graph's legend to better identify the plotted variables.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Legend**, click **Edit Properties** and select the properties you want to display in the legend.
3. Click **Apply**.
4. Click **Save**.

Adding additional statistics to a graph

Enhance a graph with summary statistics about the metrics.

Procedure

1. Click the **Report Details: Graph** tab.
2. In the **Settings for Standard Graphs** area, select the summary statistic.

Option	Description
Pinpoint	Displays the minimum and maximum values.
Boxed	Displays a box in the upper left of the graph with minimum and maximum values, the average, and the 95th percentile, which is 95% of the time the value is below this amount.
Trend	Displays a superimposed trending line over the graph, which shows the prevailing direction of the metric.
Logarithmic	Plots the values on a logarithmic scale only if there is more than one metric in the graph. This applies only to simple graphs.
Composite	Displays four supplementary mini-graphs at the bottom showing the evolution of the metric over four different time periods. The default time periods are: the last year, the last six months, the last month, and the last day. These can be configured by your system administrator.

3. Click **Save**.

Graph report parameters

You define the contextual settings of graph reports on the **Report Details: Graph** tab. This includes standard and aggregated graph reports, as well as the baseline report.

Main Y-Axis Boundaries

You can set up y-axis fixed bounds for the report to automatically focus on a relevant data range. This is useful when the acceptable values of the range are already known, such as availability data.

Main Y-Axis Scaling Mode

You can scale the values on the y-axis. If the values are too high or too low to be meaningful, you can multiply or divide by a factor that produces a meaningful number. You can also choose by-unit and select the required conversion using the offered choices. The Binary scale, such as KiB and MiB, uses 1024 as the scaling factor, and is typically used for computer memory. The Decimal scale, such as KB and MB, uses 1000, and is typically used in storage. Classic was created for the old interpretation of the multipliers where $K = 1024$.

Thresholds Definition

Critical and major thresholds appear in the graph only if you give them values. You can add as many thresholds as necessary, giving each a different value and color. If there is a second y-axis in the graph, the thresholds you set are applicable only to the main y-axis.

Thresholds have a different meaning for Time Ratio Bar Graphs. For these graphs, thresholds define the intervals for computing the time ratios and the colors of the bars. For example, two thresholds actually define three intervals. The first one gathers all values below or equal to the lowest threshold. The second one retrieves all the values above the first and below or equal to the second. The last one gathers all the values above the second threshold. The unit marked next to the input box tells you how to enter the threshold definitions

Match metrics using a

Selects the metrics to customize.

property value	Selects all metrics with a given value for a given property.
formula result	Selects the output of a formula defined on the current node or one of its children.
child report name	Selects all metrics coming from a given child report. Reports are in Browse mode.
child node name	Selects all metrics coming from a child node. If using expansion, Edit mode may generate several reports.
filter wizard	Selects metrics using a custom filter that you create using the Filter Wizard.
filter manual	Selects metrics using a custom filter that you manually create using expressions.

Configuration

This setting defines whether to display the metric against the main or the secondary axis. You can also specify the line attributes.

Properties in Legend

These properties appear in the legend to identify the plotted variables. You can add as many legends as needed and reorder the properties in the legend.

Legend Visibility

This setting indicates whether to show or hide the legend. When you hide the legend, a more compact display appears.

Legend Items

When you display only selected items, the Others section does not appear. For example, in a pie chart, the indicators in the Others section may be of no interest. Since these are not selected indicators, you can hide them by selecting display only selected items.

Summary Statistics

You have various options for displaying statistical information in a report. The default time periods are the last year, the last six months, the last month, and the last day, which are set by an administrator.

Pinpoint	Displays the minimum and maximum values.
Boxed	Displays a box in the upper left of the graph with the minimum and maximum values, the average, and the 95th percentile, which indicates whether the value is below this amount 95% of the time.
Trend	Displays a superimposed trending line over the graph showing the prevailing direction of the metric.
Logarithmic	Plots the values on a logarithmic scale if there is more than one metric in the graph. This applies only to simple graphs.
Composite	Displays four supplementary mini-graphs at the bottom that shows the evolution of the metric over four different time periods.

TopN reports

TopN reports display N number of metrics with the highest values in a specified set of data.

TopN report types

In **Report Type** on the **Report Configuration** tab, you can select these TopN reports.

Table 3 Top N reports

Type	Description
horizontal bars	Similar to a stacked chart but instead of plotting each time series, it aggregates each time series into a single value, which is shown as a percentage. This report is compatible with the single horizontal bar and the pie chart reports.
single horizontal bar	Uses the same data as the horizontal bar but displays the proportion of each child in relationship to all children in a single bar.
time ratio bars	Similar to the horizontal bar report except the bars show color pieces, each representing a certain interval value. The length of pieces show the ratio of time in percentages for the corresponding interval values. For example, you can view the percentage of time a device is experiencing critical, major, or no problems within a time range

Table 3 Top N reports (continued)

Type	Description
pie chart	Similar to the horizontal bar report except displays the results in a pie shape. This report is compatible with horizontal bars and single horizontal bar reports.

Setting the appearance of the graph metrics

You can customize the color and shape of the metrics that are displayed in the graph.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Metric Customization**, click **Add Customization**.
3. Select the metric you want to customize.
4. From the configuration options, choose whether the metric should be displayed against the main or the secondary axis, and how you want the line to be displayed.
5. Click **Save**.

Changing the legend

You can configure which properties are displayed in the graph's legend to better identify the plotted variables.

Procedure

1. Click the **Report Details: Graph** tab.
2. Under **Legend**, click **Edit Properties** and select the properties you want to display in the legend.
3. Click **Apply**.
4. Click **Save**.

Configuring the TopN sections

TopN charts retrieve the same data as normal reports, but instead of plotting a time series, they aggregate each series into a single value. Values that are too small are grouped into an Other section displayed in light gray.

Procedure

1. Click the **Report Details: TopN Graph** tab.
2. Under **Settings for TopN Graphs**, click **Section order** to select how sections will be displayed.
3. Click **Section grouping mode** and select the layout for the sections.
4. In **Number of sections**, enter the maximum number of sections to display.
5. In **Others section color**, select the color for the Other section.
6. Click **Expand Others section** to expand the Other section into subsections.
7. Click **Save**.

Selecting tooltip information

You can select the properties that will be displayed in the tooltip.

Procedure

1. Click the **Report Details: TopN Graph** tab.
2. Under **Tooltip Information**, click **Add Property**.
3. Use the Selection Property Helper to select the properties you want to add to the tooltip.
4. Click **Save**.

TopN report parameters

You define how you want data to appear in the TopN graph reports on the **Report Details: TopN** tab.

Scaling Mode

If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.

Match metrics using a

Selects the metrics to customize.

property value	Selects all metrics with a given value for a given property.
formula result	Selects the output of a formula defined on the current node or one of its children.
child report name	Selects all metrics coming from a given child report. Reports are in Browse mode.
child node name	Selects all metrics coming from a child node. If using expansion, Edit mode may generate several reports.
filter wizard	Selects metrics using a custom filter that you create using the Filter Wizard.
filter manual	Selects metrics using a custom filter that you manually create using expressions.

Configuration

Sets the color for the metric.

Properties in Legend

These properties display in the graph's legend to identify the plotted variables. You can add as many legends as you want and reorder the properties in the legend.

Legend Visibility

Controls the default visibility of the graph's legend. Hiding the legend leads to a more compact display.

Legend Items

Displaying only selected items, hides the Others section for applicable reports in the graph legend. For example, in a pie chart, the indicators in the Others section may be of no interest. Since these are not selected indicators, you can hide them by selecting display only selected items.

Section Order

Defines how to order the sections in the TopN graph.

Section Grouping Mode

Sets the layout of sections that belong to the same group. This setting is only applicable to horizontal bar graphs.

Number of Sections

Sets the maximum number of sections to display on a TopN graph. If there are more sections, they are placed in a section called Others.

Others Section Color

Sets the color of the Others section if defined.

Expand Others Section

When selected, the Others section expands into sub-sections. Otherwise, it is represented by a single aggregated section. This setting is only applicable to horizontal bars with children generated from a multiple expansion.

Map reports

A map report is basically a standard table report where each child node displays as a point in a geographical map. You can edit the location and the name of the location as well as the size, color, and shape of the icon displayed on the map.

In **Report Type** on the **Report Configuration** tab, you select the map report from the miscellaneous section.

You can edit the following attributes on a map report:

- ◆ **Location** locates an element on the map.
- ◆ **Name** provides a name to a location, such as device name, region, and sitename.
- ◆ **Size** sets the size of the icon. The size of the icon is based on a given metric value, such as the number of devices at the site.
- ◆ **Color** sets the color of the icon. The color relies on the threshold definition where green represents no thresholds, yellow represents a major threshold, and red represents a critical threshold.
- ◆ **Icon** sets the shape of the icon, such as circle, square, triangle, inverted triangle, and pentagon.

Overriding the location

You can override the location in the selected data if the location is not correctly displayed on the map. For example, "Rome Office" is displayed on the map in Rome, Ohio instead of in Italy. To fix this, you can override the location and set it to "Roma, Italia."

Procedure

1. Click the **Report Details: Map** tab.
2. Click the **Location** tab.
3. Click **Advanced Properties**.
4. In **Location overrides**, enter the new location information.

The locations entered here will have precedence over the ones that are found in the selected data. The correct format is: name=location, one pair per line.

5. Click **Save**.

Setting the icon size for child nodes

Icon size is calculated as a ratio between the node's value and the minimum and maximum value of the other nodes. The default minimum and maximum values can be changed if the theoretical min and max are known, such as for a percentage.

Procedure

1. Click the **Report Details: Map** tab.
2. Click the **Size** tab.
3. Enter the minimum and maximum values.
4. Click **Save**.

Suppose Device 1 has a CurrentUtilization of 10 percent while Device 2 has a CurrentUtilization of 13 percent. If you leave the min and max boxes empty, the maximum value is 13 and the minimum value is 10. Therefore, Device 2 will have an icon that is the maximum size (64x64 pixels) and Device 1 will have the smallest icon possible (16x16 pixels). However, as these are a percentage, you can set the maximum to 100 and the minimum to 0 to ensure that the icons for Device 1 and Device 2 are almost the same size.

Map report parameters

The settings on the **Report Details: Map** tab combine table report settings and map-specific settings to create a map report.

You set map attributes using the tabs in the **Map Item Configuration** section. You define the information to appear on the map in the **Displayed Information** section. This information uses the same column attributes, properties, and values as the table report. Any changes made to the columns in the **Displayed Information** section are replicated in the columns in the **Table Report: Details** tab and vice-versa.

Map Item Configuration tabs

Name	Defines the column that contains the icon's name. When hovering over an icon, this name is displayed in a tooltip. This information is also used as the title when opening a popup by clicking on an icon. The selected column can be a property column or an attribute column. You can select from those already configured or create a new property or attribute column.
Location	Defines the column that contains the geographical information for each icon. This column can be either a property column or an attribute column. You can select from those already configured or create a new property or attribute column. You can also set a secondary location column in case the first one does not contain information or to manually enter locations for child nodes. The format to manually enter locations is: node-name=location.
Color	Defines the column that determines the color of the icon for each child node. It is a value column and you must set a threshold on the column's value. The defined thresholds determine the icon's color on the map: ♦ Green indicates OK. ♦ Yellow indicates a major threshold was reached. ♦ Red indicates a critical threshold was reached.
Size	Defines the column that determines the size of the icon for each child node. This column can be of any type as long as it contains a numerical value. You

	can use attribute, properties, or values that already exist, or you can create a new one. The size is calculated as a ratio between the node's value and the minimum and maximum value of the other nodes. These minimum and maximum values can be overwritten if the theoretical min and max are known, such as for a percentage. To override them, simply enter values in the Minimum and Maximum boxes. For example: Assume you are mapping two devices for which the CurrentUtilization metric determines the size. Device-1 has a CurrentUtilization of 10 percent, while Device-2 has a CurrentUtilization of 13 percent. If the min and max boxes are empty, the maximum value is 13 and the minimum value is 10. Therefore, Device-2 has an icon with the maximum size (64x64 pixels) while Device-1 has the smallest icon possible (16x16 pixels). However, since these values can use a percentage, set the maximum to 100 and the minimum to 0 to ensure the icons for Device-1 and Device-2 are almost the same size, showing a realistic representation.
Search	Defines the properties that generate related reports for every icon. Adding properties in this box creates a link in each popup on the map. Clicking on that link displays a search results page with reports similar to the one on the selected node, automatically inserting the selected properties into the search query. The search results page is the same as the one that appears when entering a search query in the search text box in the upper right corner.
Icon	Defines the icon for each point. Current choices are: circle, square, triangle, inverted triangle, and pentagon. When an icon on the map is clicked, a popup displays the information contained in each defined column. The selected name column is shown as the title, and the selected location column is displayed right under the name in italics. Under the location, the other defined columns, including the ones that were used for color and size calculations, are displayed as general information in the order in which they are defined. At the bottom of the popup are links to browse the report and to search for related reports.

Displayed Information

For information on these settings, see [Table report parameters on page 34](#).

Heat map reports

A heat map report displays metrics in colors in a two-dimensional grid. You group metrics together and set the thresholds to color code the range of values for hot and cold areas. For example, red can indicate a severe problem and a lighter shade of red can indicate a warning.

You can group any combination of metric properties to define the axes of the grid. You can aggregate these metrics over time for one or both axes. The produced heat map report displays the aggregated values according to the specified time unit, which can be an hour, day, week, month, or year.

In **Report Type** on the **Report Configuration** tab, you select the heat map report from the aggregated section.

Heat Map report parameters

You set how you want information presented in a heat map on the **Report Details: Heat Map** tab.

Scaling Mode

If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.

Order

Defines how to order and group the values in the heat map.

Show Grid

When selected, draws a border around the heat map cells.

Use Gradient Color

When selected, inserts gradient colors between the heat map cells that have the closest thresholds.

Thresholds Definition

Critical and major thresholds do not appear in the report unless you give them values. You can add as many thresholds as necessary, giving each a different value and color. If there is a second y-axis in the graph, the thresholds you set here are applicable only to the main y-axis.

Properties in Legend

These properties display in the graph's legend to identify the plotted variables. You can add as many legends as you want and reorder the properties in the legend.

Treemap reports

A treemap report displays metrics in a hierarchy using size and color. Based on the tree hierarchy, each branch is given a rectangle, which is then tiled with smaller rectangles to represent the metric aggregation from each child node.

In **Report Type** on the **Report Configuration** tab, you select the treemap report from the aggregated section.

Treemap report parameters

You define how to present information on a treemap report on the **Report Details: Treemap** tab.

Number of Levels

Sets the depth of the report. You can select up to three levels. For example, to create a global report showing the total capacity of arrays, pools, and LUNS, set this level to three. To report on the total capacity for arrays and pools only, set this level to two.

Size

These settings determine the area size of the cells for the lower nodes.

Column Name	A unique name that appears in the table header.
Selected Value	Retrieves a value from the following: ♦ filter allows the configuration of a specific filter, which selects the metrics to display from the child.

	♦ property value lets you specify a property. A property field appears when you select this option. ♦ child count displays the number of child nodes for the parent node.
Filter to Apply	Selects which metrics to display in this column. This filter applies to each line of the table, which represents each child node.
Sampling Type	Selects one of the available aggregated values stored in the database. You can choose inherited from report, average, min, max, sum, last value, count, or last timestamp.
Recover...	Determines if all values are retrieved over the report time range, or only the last value. If all values are retrieved, then they are aggregated to a single value using the selected function.
Time Threshold	The time interval relative to the report time value, which acts like a tolerance interval for data retrieval. If you select real-time, the threshold is in seconds. If you select an aggregate sampling type, the time threshold is a number of periods. Note If you set zero for a time threshold while recovering only the last value using a real-time sampling period, an empty table cell may result. Always set a meaningful threshold according to the sampling period.
Scaling Mode	If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.

Color

These settings determine the color of the cells for the lower level nodes.

Column Name	A unique name that appears in the table header.
Selected Value	Retrieves a value from the following: ♦ filter allows the configuration of a specific filter through the Filter Wizard, which selects the metrics to display from the child. ♦ formula result displays the result from a formula applied to the child.
Filter to Apply	Selects which metrics to display in this column. This filter applies to each line of the table, which represents each child node.
Sampling Type	Selects one of the available aggregated values stored in the database. You can choose average, min, max, sum, the last value, count, the number of received values, or last timestamp, and the timestamp of the last received value, on the previously selected aggregate period.
Recover...	Determines if all values are retrieved over the report time range, or only the last value. If all values are retrieved, then they are aggregated into a single value using the selected function.
Time Threshold	The time interval relative to the report time value, which acts like a tolerance interval for data retrieval.

	If you select real-time, the threshold is in seconds. If you select an aggregate sampling type, the time threshold is a number of periods. Note If you set zero for a time threshold while recovering only the last value using a real-time sampling period, an empty table cell may result. Always set a meaningful threshold according to the sampling period.
Scaling Mode	If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.
Critical Level	The value for the critical threshold level. You can leave this empty to deactivate this level.
Major Level	The value for the major threshold level. You can leave this empty to deactivate this level.
Is critically ascending?	If one of the levels is empty or both levels have the same value, select or clear that setting according to your requirements. This option indicates whether the criticality is ascending where a higher number is more critical, or descending where a lower number is more critical. If a different number is given for both levels, the direction of criticality is obvious and you do not need select or clear this option.

Editing external reports

An external report is a web page that is displayed within the report pane.

Procedure

1. Click the **Report Details: External** tab.
2. Edit the **External report address** URL.
3. Click **Save**.

External report parameters

You define the external report address URL on the **Report Details: External** tab.

External report address

This is the URL of the page to display in the report. You can use `${property}` to add property values to the URL and `@{attribute}` to add report attributes to the URL.

For example, to query Bing, you can use: `http://www.bing.com/search?q=${parttype}`

The external report uses the parameters of the report's edit settings, the filter in this case, to inject properties into the external web page.

Topology reports

A topology report displays a graphical representation of the devices within a network and how they are linked together.

You can build and display transactions graphically in a topology report. You can use this report to represent a specific section of a network and to display various data between

the different elements within a section. This enables you to represent different pieces of equipment and services and the various metrics and key performance indicators linking them together.

In **Report Type** on the **Report Configuration** tab, you select the topology report from the miscellaneous section.

Report elements

Nodes and edges define a topology report. The nodes represent the different elements of a transaction or the various elements of a topology representing a network section. The edges define the relationships between the nodes. A node can have either none, one, or multiple edges linking to other nodes (n-to-n relationship). The edges are capped with an arrow, pointing from one node to another to represent the flow.

All elements in the report are plotted on a surface called a canvas. The canvas has a rectangular shape and consists of tiles called cells. A node can only be one cell and a cell can only have one node.

Topology report parameters

You define the source and layout of the report on the **Report Details: Topology** tab. An administrator must configure topology reports before you can set the source and layout for a report.

Mode

The mode of the topology report determines the source of the report. It is fully customizable.

Simple	Displays the instances of classes from the selected topology that meet the report filter. For example, if you select Router as the class, and the instances in the class are R1, R2, and R3 and the report filter is devtype=="Router" & device=="R1"; only the R1 router appears in the report.
Advanced	Use the Filter setting to restrict nodes to the report filters. This does not apply to the root nodes because they are the starting point of the graph. The root nodes are always restricted to the report filter. For example, if you use this configuration and the report filter is devtype=="Switch", and you do not click the Filter check box, the hosts appear in the report, even though they do not meet the report filter. All switches, determined by the report filter and linked to the hosts in the topology, appear in the report.
Drilldown	Renders a report based on the sub-reports of the current report. If the current report has three sub-reports, then three nodes appear. The node is based on the sub-report in Browse mode. Each node is clickable for drill down to the sub-reports. For drawing the nodes, you need a template containing the images references. If there isn't a template or if the template does not contain references to the sub-report, the interface uses the sub-reports icon. Drilldown mode displays reports that are one level below the topology report to prohibit edges from displaying.
Custom	The topology custom mode report is configured in XML files. You must place the XML file in the <code>topomap</code> directory of the <code>WebApps-Resources</code> module. For example, <code><APG>/Custom/WebApps-Resources/Default/topomap</code> .

Layout

Use this setting to determine the layout of the report.

Organic	Positions the nodes based on a physics simulation of the interacting forces.
Radial map	Positions the nodes on a tree of circles in a progressively increasing radius.
Balloon	Positions the child nodes radially around their parents.
Circular	Positions the nodes in a circle.
Force-based	Creates a force-based placement of the nodes so that all the edges are of more or less equal length and there are as few crossing edges as possible.
Hierarchical left to right	Positions the nodes in a hierarchical structure from left to right.
Hierarchical top to bottom	Positions the nodes in a hierarchical structure from top to bottom.
Tree map	Positions the nodes in a tree structure with the parents at the top right.
Manual	This setting is only available in custom mode.
Orthogonal	Positions the nodes at right angles.

Template

Use this setting to determine how the nodes and edges are drawn. For the simple and advanced modes, the `id-ref` of each node is the same as the class. For drilldown mode, the `id-ref` is the sub-report in Edit mode. For example, if you have a sub-report named "Router" and you perform an expansion on it, then each node of this expansion has the `id-ref` set to "Router."

If you select more than one template and the same node is located in multiple templates, only one is valid.

For the edge part of the simple and advanced mode, the `id-ref` is formatted as `Source.Relation.Destination`.

Overlay reports

Overlay reports define additional layers to a graph report where events are plotted. You define the events to appear in a graph in a table report.

An overlay report is a mixed report because it displays the child graph report and the child table report on the same page.

In **Report Type** on the **Report Configuration** tab, you select the overlay report from the mixed section.

Configuration process

To configure an overlay report, do the following:

1. Add an overlay report node to the tree and configure it.
2. Configure a graph report as its child.
3. Configure a table report as its child, setting the events to appear in the graph.

Overlay report parameters

You define the selection filter, the tooltip information, and the event attributes on the **Report Details: Overlay** tab.

Display Title

The title does not appear in the reports but appears in the title of the overlay configuration box.

Selection Filter

Filters events from the selected ones in the child table to appear in the overlay. This filter only needs to contain additional properties because the table performs most of the filtering.

Boundaries

Defines the zone in the graph where the events appear. It shows boundaries as percentages where zero is the bottom of the graph and 100 is the top. Use negative numbers to position the event zone below the graph. Use the slider for adjustments. Your selection appears with the overlay title at the top of its configuration box.

Number of Lines

The number of lines to show the events in the graph. You can use one line per 10% of height. The result is alternating lines distributed vertically to plot the events.

Main Color

The color value to use for the background of the whole zone. Set to white, `#ffffff`, to disable it.

Display Level

Events and time series are rendered on top of one another. This setting determines how to arrange events and time series in an overlay. To draw overlays under the graph lines, use a negative value. To order different overlays, use numbers for the respective display levels to arrange them from the lowest value to the highest value. You can also use top or bottom as shortcuts for min and max. If multiple overlays overlap, you can choose the order in which they appear, as they are drawn from the lowest value to the highest.

Display Mode

The way to render events. Depending on your selection, settings for events with duration are toggled to display those applicable to lines or areas.

Event Limit

The maximum number of events displayed in the overlay. Events are displayed using the table order, and after the limit is reached, the rest of the events are not displayed.

Tooltip Title

The name of the property whose value is the tooltip title.

Events Color

The HTML color code for momentary events.

Point Size

The size of the points for momentary events.

Marker Type

The type of marker for momentary events.

Events Color

The HTML color code for durable events.

Line Width

The line width for events with duration.

Line Style

The line style for events with duration.

Start Marker

The marker for the start of events with duration.

End Marker

The end marker for events with duration.

Status reports

A status report aggregates data for each child node using a selected function over a time period. Each child node displays as a color symbol that can show its aggregated value and node name. This enables you to quickly view the status of nodes in relationship to one another, such as their average CPU utilization.

You select the size and symbol to use for the child nodes as well as the thresholds and their colors.

In **Report Type** on the **Report Configuration** tab, you select the status report from the aggregated section.

Status report parameters

You set how you want information presented in a status report on the **Report Details: Status** tab.

Marker Type

Sets the size and symbol to use when reporting on the child nodes. The size values are:

Nano	Generates a high-density color-coded status report that is similar to a heat map. It shows the aggregated value and node name in a tooltip. This is the smallest size for a symbol.
Micro	Generates a high-density color-coded status report that is similar to a heat map. It shows the aggregated value and node name in a tooltip.
Small	Generates a color-coded status report with small status symbols, each displaying its aggregated value and node name.
Medium	Generates a color-coded status report with medium status symbols, each displaying its aggregated value and node name.
Large	Generates a color-coded status report with large status symbols, each displaying its aggregated value and node name.

Show Value

Indicates whether to display the aggregated value inside the symbol and whether to round this value.

Scaling Mode

If the unit of the values does not correspond to what you want to display, select multiply or divide from this list. In the corresponding field, you can then enter the factor to apply, which allows you to change from bytes to kilobytes, for example.

Order

Indicates whether to show the child nodes in ascending or descending order.

Top N Mode

Defines whether to display all the status symbols in a report or a certain number that you specify here.

Thresholds Definition

Critical and major thresholds do not appear in the report unless you give them values. You can add as many thresholds as necessary, giving each a different value and color. If there is a second y-axis in the graph, the thresholds you set here are applicable only to the main y-axis.

Formula configurations

You can apply formulas to any node of the tree to perform advanced computations on any number of values. All formulas have at least one result, and each result has additional options for displaying them in different contexts.

Where you use formulas depends on the type of report and the hierarchy of its nodes. When applying a formula to a node, you can plot its result on a graph, insert it into another formula, pass it to another report node, display it in child-based parent report, and display it in a value column in a table on a parent node.

Results can vary when applying formulas to nodes that are subject to an expansion. Depending on the formula type, the aggregation of input parameters from child nodes in a formula can apply to their parent node. You can then view the results in the parent node as an aggregation of the results returned by the child nodes.

By using the report editing tools, you can copy formulas from one report to another report.

Handling formula results

The way you apply formulas depends on the how you want to process the results:

- ◆ If you want to plot the results of a formula on a simple chart or bar chart for a node, you can apply the formula to the node.
- ◆ If you want to use the results of a formula in a table report, apply the formula to a child of the parent table report node.
- ◆ You can use formula results from the current node or direct children as parameter inputs on the current node. This lets you can add several formulas to a node whose results are combined.
- ◆ If you want to use the results of a formula in any child-based report, such as a child-based chart or stacked bars, apply the formula to a child of the node of one of these report types.
- ◆ If you want to pass the results of a child node to a parent of the current node, you can apply a formula to the current node. You can apply formulas to nodes in a hierarchy to pass the results further up the tree until you can access them at the desired node level.

Formula parameters

You can add pre-defined formulas on the **Formula** tab and then customize them according to your reporting requirements.

Parameter

The parameter values can be over an interval of time defined in the formula, such as the results of a filter or a child formula. Not all formulas have parameters to enter.

Filter on this Node	Uses a standard filter to select the parameter to use. Values are subject to the filter of the current node, its parent nodes, and any other filters applied to the user or role. You can select any filter, but there can be specific values that are required.
Formula Result	Uses the result from a formula as the parameter. A list appears that displays the available results you can use. ◆ On a child node: results from formulas of nodes that are direct children of the current node. ◆ On the current node: results from other formulas on the current node, which can be used as a parameter value for this formula. Result names are taken from the default name the formula gives the result or the name you give the result.
Constant Value	Enters a float value to use as a constant input such as 3.0.
Property Value	Click the Property Selection Helper icon and select a property whose value is the input parameter. You can select only one property to enter. You can also type a property name in the Property field. Property values either register as a numerical value, or if not possible, do not return a value.
Combined Parameters	Enables you to combine parameter instances and types of Filter on this Node, Formula Result, Constant Value or Property Value. How the combined parameters are treated depends on the formula. Usually they are aggregated spatially using the inherited report duration and aggregation settings, or by the time settings defined in the formula. Some formulas treat combined parameters in other ways. Refer to the formula and parameter descriptions for more information.
Empty Parameter	The parameter has no value (null).
Attach properties	Enables you to attach the specified properties from the node to a parameter. If this parameter has a filter, it is also applied to the properties.

Setting

These are constant values used by the formula. You can change the default settings. These can determine, for example, the time range of the report, values that determine SLA objectives and settings that affect the aggregation the formula uses. Not all formulas have settings.

Result

The results that are returned by the formula. Each result has a default name that you can edit. The names are used for these purposes:

- ◆ In the legend of graphs to identify each result if Show in Graphs is enabled for them.
- ◆ To identify the result when using it as an input parameter for another formula, and when displaying it in a value column of a parent node table report.

Names must be unique. If there are two results with the same name and you want to display the results in a parent table report, only one instance of the name is available, and the result displayed is the aggregate of the formula results with this name. Result names are case sensitive so you can change the capitalization to differentiate results.

Show in graphs	For simple chart and bar chart graphs, if you select show in graphs for a formula result, the metrics of the node are not displayed. Only the formula results you enabled show.
Default results	You can choose one result from the results for all of the formulas on the node to display in a direct parent report that is a child-based report. To do this, select default result for the result. If the result is applied on an expansion, the result plotted in the parent node is an aggregation of all child results.

You can pass a node's result to other nodes higher in the hierarchy to use in other formulas. The child result of the current node is passed to the parent. You can also pass a result to another formula on the same node.

Creating pre-generated reports

Pre-generated reports are reports that are automatically generated on a predetermined schedule. Pre-generated reports can save users time when they are used to generate reports that take a long time to process. Only administrators can create pre-generated reports.

Procedure

1. In Edit mode, click the **Pre-Generation** tab.
2. In **Schedule this report for**, select the user or role for which the report will be generated.

The assigned role requires at least one internal user: if the role has only externally authenticated users or no users at all, the report will be not generated.
3. In **Name**, enter a name for the report.
4. Use the **Schedule** options to select the time to generate the report.
5. Select the instances of this report to generate.

For example, if a report node has been copied and placed somewhere else in the tree, this menu will allow both of those nodes to be generated.
6. Check to send an email about the report generation to the selected recipient.

Report URL syntax

The reports in the interface are available to third-party client tools that use the URL syntax and the auto-login functionality.

URLs

The base URL is: `http://[APGserverIP:port#]/APG/?param=value¶m=value`

The URL to display reports on a full page is: `http://[APGserverIP:port#]/APG/report.jsp?param=value¶m=value`

The URL to display reports in jpg, png, pdf, svg, xls, and csv formats is: `http://[APGserverIP:port#]/APG/report.format?param=value¶m=value`

You must replace *format* in the above URL with the desired format. For example: `http://[APGserverIP:port#]/APG/report.csv?param=value¶m=value`

Tree management

Syntax	Description
<code>select=<nodeid></code>	Displays the report of the specified node.
<code>expand=<nodeid></code>	Expands the selected node and then the parent nodes.
<code>collapase=<nodeid></code>	Collapses the selected node and any child nodes.
<code>collapaseALL</code>	Collapses the whole tree.

Report settings

Syntax	Description
<code>display=<code></code>	The code numbers for the display modes: ◆ 0=normal display mode ◆ 1=summary display mode with one graph per line ◆ 2=summary display mode with two graphs per line ◆ 3=summary display mode with three graphs per line
<code>node=<code></code>	The codes for the report modes: ◆ srt=table report mode ◆ lrt=leaf table report mode ◆ frg=graph report mode ◆ srg=children report mode ◆ nrx=node report mode ◆ vrx=baseline report mode ◆ stk=stacked chart report mode ◆ hb=horizontal bar chart report mode ◆ pie=pie chart report mode ◆ gg=gauge chart report mode ◆ ico=icon report mode ◆ mix=basic mixed report mode ◆ dmx=mixed default preferences report mode
<code>period=<seconds></code>	The period, in seconds, of the selected aggregate. Zero indicates real time. This period must exist in the database.
<code>type=<code></code>	Not applicable when the period is zero. The code numbers for the aggregate types are: ◆ 3=average ◆ 4=minimum ◆ 5=maximum

	◆ 6=sum ◆ 7=last value ◆ 8=number of aggregated values ◆ 9=timestamp of the last aggregated value
<code>var_idx=<id>_<id>_<id>...</code>	A list of indexes separated by underscores that restricts the displayed variables.

Graph display preferences

Syntax	Description
<code>width=<pixels></code>	The graph width in pixels.
<code>lower=<value></code>	The lowest value displayed on graphs.
<code>upper=<value></code>	The highest value displayed on graphs.

Time management

Syntax	Description
<code>durationType=<code></code>	The code to use as the default for the report time range description: ◆ n=not applicable because duration is based on something else such as a timestamp ◆ p=previous duration ◆ l=last duration ◆ c=current duration
<code>duration=<code></code>	The code to use as the default for the report time range: ◆ a=custom code that is a combination of numbers and units ◆ s=seconds ◆ m=minutes ◆ h=hour ◆ d=day ◆ w=week ◆ M=month ◆ y=year You can use as many duration codes as needed for the report time range but you cannot duplicate a code. You can also use calendar in the duration code followed by

	start and end or start_ts and end_ts, as explained in the following rows.
start=<date>end=<date>	The time range of the requested report. The <date> field has the YYYY-MM-DD HH:MM format. Note URL encoding may transform spaces into plus signs and colons and periods into %3A.
start_ts=<timestamp>end_ts=<timestamp>	The time range using timestamps of the requested report. These are UNIX timestamps.

Advanced time management

Syntax	Description
itz=<timezone code>	The time zone for the report, which uses standard zone names such as America/Montreal.
tf=<time filter.expression>	The maintenance period for the report. The <time filter.expression> field consists of <dom>;<dow>;<hod> to represent the following: ♦ dom=comma-separated list of days in a month between 1 and 31. For example, 1, 2, 10 represent the first, second, and tenth day of the month. ♦ dow=comma-separated list of days in a week between 1 and 7, with 1 starting on Sunday. For example, 1 and 2 represent Sunday and Monday. ♦ hod=comma-separated list of hours in a day between 0 and 23. For example, 12 and 13 represent 12pm and 1pm.

Query and variable selection

When using this syntax, always check that the URL encoding did not change spaces into plus signs.

Syntax	Description
search-base=<base>	The node from which the search begins. The <base> field can contain the root node, which is the default, or another node combined with a parameter.
q=<search string>	The query string for the standard or quick search mode. It is a space-separated list of tokens.
qsp=<search properties>	The space-separated list of properties in which to search.

<code>qsg=<grouping properties></code>	The space-separated list of node expansion to group query results.
<code>qf=<search filter></code>	The variable selection filter.
<code>qg=<grouping properties></code>	The space-separated list of node expansion to group query results.

Auto-login

Syntax	Description
<code>autologin=<user>:<pass></code>	Bypasses the Login page by supplying user credentials to the URL. Since the colon is part of the auto-login syntax, the username and password cannot contain a colon.

CHAPTER 3

Managing ReportPacks

This chapter includes the following topics:

◆ Creating a ReportPack.....	64
◆ Uploading a <ReportPack>.arp file.....	64
◆ Exporting a ReportPack.....	64
◆ Deleting a ReportPack.....	65
◆ Adding a new report template to a ReportPack.....	65
◆ Copying a report template.....	65
◆ Deleting a report template from a ReportPack.....	66
◆ Generating a report template in XML format.....	66

Creating a ReportPack

A ReportPack is a container for the report templates that you create. A new ReportPack is created as a root node in the tree. Each report template is represented as a report in the interface.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks > New ReportPack**.
3. Type the **Name**, **Version**, and **Description** of the ReportPack.
4. Click **Save**.

Uploading a *«ReportPack».arp* file

You can add a *«ReportPack».arp* file that contains required information and report templates. The ReportPack appears as a root node branch in the tree.

Procedure

1. Click **Administration**.
2. Click **ReportPacks**.
3. Click **Upload a ReportPack** and select **Choose File** to browse to the *«ReportPack».arp* file.

Upload only .arp file types.

4. Click **OK**.

Results

A message displays,

The file has been uploaded successfully.

Click **Continue**.

Exporting a ReportPack

You can export a ReportPack to back up the *«ReportPack».arp* file or for use in another system.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Click **Export** to download the *«ReportPack».arp* file that contains the required information and templates.
5. Click **Save** to navigate to the download location.

Deleting a ReportPack

When a ReportPack is no longer needed, you can remove it and its report templates.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Click **Delete** to remove the ReportPack and its templates.
5. Click **Ok** at the warning message.

Adding a new report template to a ReportPack

You can add a new report template to one or more nodes in the tree. Each report template represents a report in the interface.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Click **New Template**.
5. Type the **Name**, **Version**, and **Description** of the report.
6. To add this report to a specific node in the report tree, click **Hook to Another Node** and select where to place the report.

You can add this report to as many nodes as you want by clicking **Hook to Another Node**. If nothing is selected, the template is added as one of the root nodes of the tree.

7. Click **Save**.

Copying a report template

You can copy a report template to another node in the tree.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Select the report to copy.
5. Click **Move / Copy** and select where to copy the report.

Deleting a report template from a ReportPack

If a report is no longer needed by a ReportPack, you can remove it.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Select the report to delete.
5. Click **Delete**.
6. Click **Ok** at the warning message.

Generating a report template in XML format

You can generate a report in XML format to store on a computer.

Procedure

1. Click **Administration**.
2. Under **ReportPacks**, click **List / Import ReportPacks**.
3. Select the ReportPack.
4. Select the report template.
5. Click **Export Reports**.
6. Click **Save** and navigate to the location to place the report.
7. Click **Save**.

CHAPTER 4

Managing Users

This chapter includes the following topics:

◆ How the interface manages user rights	68
◆ Adding a user	68
◆ Setting the access rights of a user	69
◆ Setting the viewing rights of a user	70
◆ Testing user settings	70
◆ Copying users	70
◆ Editing users	71
◆ Deleting users	71
◆ Viewing users	71
◆ Creating a role	71
◆ Adding members to a role	72
◆ Adding external members to a role	72
◆ Establishing inheritance between roles	72
◆ Setting role access to report templates	73
◆ Setting component access and restrictions	73
◆ Assigning administration tasks to a role	74
◆ Editing a role	74
◆ Deleting a role	75
◆ Removing a user from a role	75
◆ Resetting default role settings	75
◆ Viewing roles	75
◆ Creating a profile	76
◆ Customizing a profile	76
◆ Adding members to a profile	76
◆ Adding external members to a profile	77
◆ Editing profiles	77
◆ Removing a member from a profile	77
◆ Viewing profiles	78

How the interface manages user rights

A global administrator can assign access rights to users and groups to control what they can see and do in the interface.

Roles

Roles group users together to help you manage access to reports, modules, and functionality, such as report editing and report searching. You can also assign administrative tasks to a role.

You assign a user to a role when creating or modifying a user account or a role. There are several default roles that you can use and you can create your own roles. You can assign a user to one or more roles.

If you do not assign a user to a role, the user only sees the Scheduled Reports, Stored Reports, and Favorite Reports branches in the report tree.

Inheritance of privileges

You can define a parent-child relationship in which one role inherits the privileges of another role. In this relationship, the child role uses the parent access rights and restrictions for viewing reports and components and using functionality.

Master filter

A master file provides you an additional layer to manage access to metrics shown on reports. You can set up a master filter to restrict the users of a role from viewing certain metrics on reports.

You can also define a master filter for a particular user, which is combined with the group's master filter to place more viewing restrictions on that user.

Profiles

A user contains personal information, such as login credentials. You associate a user with a profile to define the global characteristics of a user, such as language and time zone. Users are automatically assigned to the default profile but you can add profiles according to your company's global reporting requirements.

Advanced mode

You perform most user management functions in standard mode. However, to establish inheritance between roles, assign administration tasks to a role, and reset default roles to their original settings, you must use advanced mode.

Adding a user

When you add a user, you set the login credentials and email address in which to send notifications about reports.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Click **New User**.
4. On the **User Data** tab, type the **User Login**.

This value is case sensitive.

5. For the password, do one of the following based on the type of authentication:

Option	Description
Internal authentication	Type the user login password.
External authentication	If using an external authentication mechanism, such as LDAP, do not type the user login password since it is not needed.

6. Select the **Title** of the user.
7. Type the **First Name** and the **Last Name** of the user.
8. Type the **Email Address**.

This address receives notifications about reports, such as when a certain report is generated.

9. Click **Save**.

A user can override these settings on the **User Settings** page.

Setting the access rights of a user

You can indicate whether a user is a global administrator with full privileges or a normal user in which you can assign access rights to reports, components, and report templates using roles.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Edit**.
4. Click the **User Status** tab.
5. For **User Status**, select one of the following:

Option	Description
Normal User	Default status for all users. You can set various access restrictions on components and reports on a normal user.
Global Administrator	Full access to all components and reports. You cannot disable a global administrator from login.

6. If you want to disable the user from logging in, click **Disabled?**.

This is useful when performing maintenance.

7. Select the **Profile** for the user.

The user is automatically added to the default profile if you did not create any profiles.

8. Use **Add Role** to assign the user to a default role or other roles that you created.

The role you assign the user determines which reports a user can see and modify. The Full Control Users default role enables a user to access and modify all available reports in the tree as well as access most modules and tools.

9. Click **Save**.

Setting the viewing rights of a user

You can create a master filter to restrict the data a user can see on a report and restrict a user from viewing custom reports.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Edit**.
4. Click the **Other Options** tab.
5. In **Master Filter**, define one or more filters that place viewing restrictions on the user.

These filters are combined to create the master filter. If the user is assigned to a role with a master filter, then the role's master filter is combined with this master filter to place more restrictions on what the user sees.

6. In **Custom Reports**, indicate whether you want the user to view the reports in the **My Reports** area.
7. Click **Save**.

Testing user settings

You can test which reports a user can see.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Test User**.

You are launched into the interface where you can view reports as the specified user.

4. Test the user by browsing and editing reports.

Any changes you make are saved.

5. When finished with testing, click **Return to Self**.

Copying users

You can copy the settings of an existing user to create a new user.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Copy**.
4. Type the settings for this new user.
5. Click **Save**.

Editing users

You can modify the settings of existing users.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Edit**.
4. Make the edits.
5. Click **Save**.

Deleting users

When you delete a user, all the user properties are removed.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.
3. Select the user and click **Delete**.
4. Click **Ok** when the warning message appears.

Viewing users

You can generate a list of existing users that shows the status and profile and whether the profile is currently enabled.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Users**.

Creating a role

Roles group users together to help you manage access to resources such as reports and metrics. You can assign a user to one or more roles.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles** > **New Roles** > **Main Properties**.
3. Type the **Name** and **Description** of the role.
4. In **Master Filter**, define one or more filters that place viewing restrictions on the users of the role.

These filters are combined to create the master filter. If a user in the role is also assigned to a master filter, then the user's master filter is combined with this role's master filter to place more restrictions on what the user sees.

5. To disable users of this role and any associated child roles, click **Disabled?**.

These users cannot log in until you clear this check box. This setting is useful when performing maintenance on a group of users.

6. Click **Save**.

Adding members to a role

You can group users with the same access requirements and restrictions into a role.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click the **Members** tab.
5. Use **Add to Role** to assign users to the role.
6. Click **Save**.

Adding external members to a role

If you are authenticating users that use an external mechanism such as LDAP, you can assign this group of users to a role.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click the **External Members** tab.
5. Type the name of the group as defined in the external authentication mechanism in **Bind an external group to this**.
6. Click **Add to this role**.
7. Click **Save**.

Establishing inheritance between roles

If you want to set up a relationship in which one role inherits the privileges of another role, create a parent-child relationship. In this relationship, the child role uses the parent access rights and restrictions for viewing reports and modules and using functionality.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. On the **Role Management** page, verify that you are working in advanced mode. If not, click **Standards Mode**.
4. Select the role.
5. On the **Main Properties** tab, select the parent role from the **Parent Role** list.
This list only appears when working in advanced mode.
6. Click **Save**.

Setting role access to report templates

For each role, you can specify access rights to individual reports, all reports in a ReportPack, or all reports in all ReportPacks.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click the **Template Access** tab.
5. You can do the following to set role access to report templates:

Option	Description
Access for individual templates	Sets the role access rights for an individual template. The system looks at this first when determining access. You can provide different access levels to a template than its ReportPack. For example, a template can have read-write access even though its ReportPack has no access.
Access for individual ReportPacks	Sets the same role access rights for all templates in a ReportPack. You can provide different access levels to the templates in a ReportPack. For example, you can set read-only access to a ReportPack and set no access to a template in that ReportPack.
Default access for all templates	Sets the same role access rights for all templates in all ReportPacks. The system looks at this last when determining access. You can set different access levels to both ReportPacks and templates. For example, the default access can be read-write but you can set no access to specific ReportPacks and templates.

6. Click **Save**.

After you finish

You can use a master filter to restrict what users of a role see on a report.

Setting component access and restrictions

You can set role access rights to various components and place role restrictions and limits on functionality within certain components. For example, you can restrict users of a role from using Edit mode in the interface. If a role is a child, you can ensure the access rights of its parent are enforced.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click the **Modules and Restrictions Access** tab.
5. Set the role access for a component:

Option	Description
Inherit	If this is a child role, Yes appears under Inherit . For a child to inherit the access rights of a parent, you must select Yes under Enforce .
Enforce	Provides role access to the component.

6. To set role restrictions and limits on specific functionality within a component, click the arrow icon in front of its name.

A component without any restrictions has a circle in front of its name.

7. Specify the restrictions as needed.
8. Click **Save**.

Assigning administration tasks to a role

You can set up a role to act as an administrator. The users of this role can perform user and role management on this role and any child roles. In a large multi-tenancy environment, this is useful because user maintenance does not affect other customers or departments in the system.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Verify that you are working in advanced mode. If not, click **Standard Mode**.
4. Click the **Role and User Management** tab.
This tab only appears when working in advanced mode.
5. Select the role and do the following:
 - To manage this role and any child roles, click **Roles**.
 - To manage the users of this role and any child roles, click **Users**.
6. Click **Save**.

Editing a role

You can modify the settings of an existing role.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Make the edits.
5. Click **Save**.

Deleting a role

If a role is no longer required, you can delete it.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click **Delete**.

Removing a user from a role

When a user no longer belongs to a role, you can remove that user from the role.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Select the role.
4. Click the **Members** tab.
5. Use **Remove from Role** to remove the user.
6. Click **Save**.

Resetting default role settings

By resetting roles to their original defaults, you are removing all modifications made to the Full Control Users, Read-Only & Real-Time Grapher Users, Read-Only Users, and Read-Write Users default roles.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.
3. Verify that you are working in advanced mode. If not, click **Standard Mode**.
4. Click the **revert roles to factory defaults** link in the text.
5. Click **Ok** when the warning message appears.

The original default roles appear on the **Role Management** page.

Viewing roles

The Role Management page gives you an at-a-glance view of the roles on your system, their descriptions, and how many users currently use each role.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Roles**.

Creating a profile

Profiles enable you to group users and roles together according to your company's global reporting requirements. You can create as many profiles as needed for your locales and time zones.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles** > **Main Properties**.
3. Type the **Name** of the profile.
4. Type the **Description** of the profile.
5. Click **Save**.

Customizing a profile

You can customize a profile to meet the language and time zone needs of its users. You can also indicate which report you want to display at user login.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.
3. Select the profile.
4. Click the **Customizable Settings** tab.
5. Select the **Locale** for the entire interface.
The default language is the language of the server.
6. Select the **Time Zone** for the profile.
7. Select an image from the logo list to display above the report tree and on the **Login**, **User Interface**, and **Administration** pages.
8. Type the location of the **Login Report**.
This report appears after users log in. If left blank, the default login report appears.
9. Type the **Auto Refresh Rate** in seconds.
10. Click **Save**.

Adding members to a profile

Users are automatically assigned to the default profile. You can add these users to new profiles as well as remove users from the default profile. You can also add roles to a profile.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.
3. Select the profile.
4. Click the **Members** tab.

5. Use **Add to Profile** to assign users and roles to the profile.
6. Click **Save**.

Adding external members to a profile

If you are authenticating users that use an external mechanism such as LDAP, you can assign this group of users to a profile.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.
3. Select the profile.
4. Click the **External Members** tab.
5. Type the name of the group as defined in the external authentication mechanism in **Bind an external group to this**.
6. Click **Add to this profile**.
7. Click **Save**.

Editing profiles

You can modify the settings of existing profiles.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.
3. Select the profile.
4. Make the edits.
5. Click **Save**.

Removing a member from a profile

When a user and role are no longer members of a profile, you can remove them from the profile.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.
3. Select the profile.
4. Click the **Members** tab.
5. Select the member to remove from **Users with this Profile**.
6. Click **Remove from Profile**.
7. Click **Save**.

Viewing profiles

The Profiles page gives you an at-a-glance view of the profiles on your system, their descriptions, and how many users currently use each profile.

Procedure

1. Click **Administration**.
2. Under **User Management**, click **Profiles**.

CHAPTER 5

Administering the Interface

This chapter includes the following topics:

◆	Setting the login properties.....	80
◆	Customizing the look of the interface.....	80
◆	Enabling keyboard shortcuts.....	80
◆	Modifying the default profile settings.....	81
◆	Editing the aliases given to time periods.....	81
◆	Disabling user logins.....	81
◆	Cancelling a user session.....	81
◆	Accessing the scheduled reporting options.....	82
◆	Editing the properties of scheduled reports.....	82
◆	Listing reports that are generating.....	83
◆	Deleting stored reports.....	83
◆	Accessing modules.....	83
◆	Deleting metrics.....	83

Setting the login properties

You can define the welcome message and the **Login** page autocomplete behavior.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **APG Portal** › **Frontend Status**.
3. You can select one of the following for the **Login** page:
 - Disable browser autocomplete for both the username and password.
 - Disable browser autocomplete for the password only.
4. Type the customized splash screen message in **Welcome Message**.
5. Click **Save**.

Customizing the look of the interface

You can change the look of the interface by applying CSS style sheets and JavaScript files to customize interface elements such as buttons and menus.

Procedure

1. Click **Administration**.
2. Under **Frontend Management** click **APG Portal** › **File Customization**.
3. Upload a CSS file.
4. Upload a Javascript file.

If you need to edit the CSS and Javascript files, you can download them for editing and then upload them.
5. Click **Save**.

Enabling keyboard shortcuts

Keyboard shortcuts enable you to perform common tasks quickly. For example, you can switch between Browse and Edit mode by pressing **Alt + e**.

Procedure

1. Click **Administration**.
2. Under **Frontend Management** click **APG Portal** › **File Customization**.
3. Check **hotkeys.js**.
4. Click **Save**.
5. Press **Alt + h** to view the keyboard shortcuts.

Modifying the default profile settings

You can also modify the language, time zone, and logo defined for the default profile on the **Global Portal Properties** page.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **APG Portal** › **Default Display**.
3. Upload a new logo.
4. Select the **Locale**.
5. Select the **Time Zone**.

Any changes made to the default profile on the **Profile Modification** page override these settings and vice versa.

6. Click **Save**.

Editing the aliases given to time periods

You can edit the default aliases that represent a time period and a language. For example, quarter is an alias for three months in English. These aliases appear in report titles.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **APG Portal** › **Duration Aliases**.
3. Edit the default values for the aliases.
4. Click **Add period alias** to add another alias.
5. Click **Save**.

Disabling user logins

You can disable all non-administrator logins to perform maintenance.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **APG Portal** › **Frontend Status**.
3. Click **In Maintenance?** to place the **User Interface** in a maintenance state.

This disables all logins for non-administrators until you clear this check box.

4. Click **Save**.

Cancelling a user session

If you need to log a user out of the interface, you can cancel that user's login session.

Procedure

1. Click **Administration**.

2. Under **Frontend Management**, click **User Sessions**.
3. Select the user and click **Kick Out**.

You cannot cancel the session of an administrator.

Accessing the scheduled reporting options

You can edit the properties of a report, launch a report immediately, cancel a running or queued report, remove a report from a scheduled list, and disable a scheduled report from running.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **Scheduled Reports**.

All scheduled reports appear, including those that are pending, currently running, pre-generated, such as the daily health report, and any invalid reports.

3. Select a report. The options that appear depend on the status of the report. For example, **Launch Now** only appears if the report is not running.

Option	Description
Edit	Modifies the properties of the report.
Enable	If the report is disabled, enables the generation of the report based on its schedule.
Disable	If the report is enabled, disables the generation of the report. The report cannot run until it is enabled.
Launch Now	Runs the report immediately.
Abort Now	Stops a report that is currently running.
Delete	Removes the report from the scheduled list.

Editing the properties of scheduled reports

You can rename a point-in-time report, modify its schedule, change the recipients of a report, and change the instances of the report to generate.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **Scheduled Reports**.
3. Select a report and click **Edit**.
4. **Schedule this report for** all the users or only the users in a specified role.
5. To rename the report, type the **Name**.
6. To generate this report at the scheduled time, click **Active**.
7. In **Schedule**, do one of the following:
 - Reset the existing settings to schedule the report.
 - Click **Advanced** to schedule the report using cron.
8. Select the instances of the report to generate in **Report Link**.

9. To send an email about the generation of the report, click **Send confirmation email**, and type the email recipients separated by commas.
10. Click **Save**.

Listing reports that are generating

At any point in time, you can determine which reports are currently running.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **Running Reports**.

Deleting stored reports

When a stored report is no longer needed, you can remove it.

Procedure

1. Click **Administration**.
2. Under **Frontend Management**, click **Stored Reports**.
3. Click the check box in the first column of the report to delete.
4. Click **Delete**.

Accessing modules

Modules are components that collect, organize, and process data about hardware components, applications, storage devices, and databases.

You install and manage modules in **Centralized Management**. Some modules require a license for installation.

Procedure

1. Click **Administration**.
2. Under **Modules**, select a module.

Deleting metrics

You can remove metrics that are no longer needed from the database.

Before you begin

If you added databases to the system, verify that the resource names and resource links of these databases are in the `APG.xml` file. This enables you to delete metrics.

Procedure

1. Click **Administration**.
2. Under **Modules**, select **Management of Database Metrics**.
3. Create the **Filter** for displaying the metrics to delete.
4. Type the **Maximum results** to appear.
5. To show the timestamp of each metric, click **Show last timestamp for all results**.
This can increase the search time.

6. Select the **Properties to show** for the search results.
7. Click **Query**.
8. When the search results appear, click each metric to delete.
9. Click **Delete**.

Until the data in the database is refreshed, these metrics are not removed and may still appear in reports.

CHAPTER 6

Customizing User Settings

This chapter includes the following topics:

◆ Modifying your user data	86
◆ Setting your reporting preferences	86
◆ Saving the report tree	86
◆ Uploading a report	87

Modifying your user data

You can change the user data that was created by a global administrator. You can modify your password, title, name, and email address in which to send report notifications.

Procedure

1. Click **Settings**.
2. On the **User Data** tab, make the changes as needed.
3. Click **Save**.

Setting your reporting preferences

You can set the default behavior for displaying reports.

Procedure

1. Click **Settings**.
2. Click the **Preferences** tab.
3. Select the **Language** to use.
4. In **Root Display**, select how you want the root nodes to display multiple reports on the report page.

A root node is the uppermost node of each branch in the tree, such as **Stored Reports** and **Favorite Reports**. The default behavior displays one report after another report.

5. In **Root Report**, select which report type to use for the root nodes.
6. Set the **Auto Refresh Rate**.

This value comes from the profile configuration.

7. In **Background Reports**, select one of the following:

Option	Description
Show the Question	When a report takes a long time to generate, a question appears asking if you want to wait until the report finishes or if you want to finish the report in the background and then save it.
No Background Reports	No question appears because generating a report in the background is not an option.

8. In **Stored Reports**, select whether you want a confirmation email sent to you after a report is saved.
9. Click **Save**.

Saving the report tree

Always make a back up copy of the entire report tree before you customize it.

Before you begin

Ensure you are logged in with User Interface mode enabled to gain access to the **Settings** tab.

Procedure

1. Click **Settings**.
2. Click the **Custom Tree** tab.
3. Click **download a backup** in the text.
4. Click **Save** and navigate to the place to store the report tree.

Uploading a report

You can add a report as a child to a node in the report tree.

Procedure

1. Click **Edit Mode**.

You can only perform this task in edit mode.

Browse Mode displays when you are in edit mode.

2. Select the node under which to add the report.

If you do not select a node, the report is added to the last root node in the tree.

3. Click **Settings**.
4. Click the **Custom Tree** tab.
5. Click **Browse** to locate the report.
6. In **Conflict Management**, select one of the following when uploading a node to a node with the same ID:

Option	Description
duplicate	Provides another ID to the new node. The existing node keeps its unique ID.
overwrite	Replaces the existing node with the new node.

7. Click **Save**.

CHAPTER 7

Licensing

This chapter includes the following topics:

- ◆ [Adding a SolutionPack license](#)..... 90

Adding a SolutionPack license

If you want to add a SolutionPack to an installed suite, create a Service Request to obtain the new license.

Some SolutionPacks are available for purchase with the suite. To add one of these licenses to an installed suite, create a Service Request to obtain a new ELMS suite license and reinstall the suite license.

Some SolutionPacks are not associated with the suite. To add one of these licenses to an installed suite, create a Service Request to obtain a license for the new SolutionPack and install the SolutionPack license.

There are two type of licenses:

1. Permanent--Permanent licenses do not expire.
2. Temporary--Temporary licenses have expiration dates. When temporary licenses expire, the module license checks fail and the modules stop working. If the frontend module fails, you cannot access the product.

If you change the configurations of a host, licenses for SolutionPacks that are not part of the suite may not work. In that case, obtain a new host ID and request a new permanent license.

Adding licenses

You must upload a license file whenever you install a new SolutionPack or ReportPack. The license enables views and reports to display.

Before you begin

- ◆ Obtain the required license file by creating a Service Request on EMC Online Support.

Procedure

1. Log in to the host with administrator credentials.
For example, `http://Frontend-hostname:58080/APG`
2. Click **Administration**.
3. In the **Centralized Management** pane, select **Licenses**.
4. In the Physical Overview tree, select **Licenses Management**.
5. In the **Licenses Listing** pane, click **Upload**.
6. In the **License Upload** dialog box, browse to your license file. Click **OK**.
Licenses in a .zip file are extracted and appear in the **Licenses Listing** pane.
7. When **License Upload Complete** appears, click **OK**, then click **Continue**.
8. If you have remote servers in a distributed environment, click **Synchronize** to add the licenses to the remote servers.

Deleting licenses from core software

You can select and delete specific licenses or delete all existing licenses from the core software server.

Procedure

1. Log in to with administrator credentials.

For example, `http://Frontend-hostname:58080/APG`

2. Click **Administration**.
3. In the **Centralized Management** pane, select **Licenses**.
4. In the Physical Overview tree, select **Licenses Management**.
5. In the **Licenses Listing** pane, select one or more license checkboxes.
6. Click **Delete**.
7. If you have remote servers in a distributed environment, click **Synchronize** to remove the licenses from the remote servers.

Verifying core software licenses from Centralized Management

Verify whether core software licenses are valid or expired for all servers in your environment from Centralized Management. The local and remote servers listed in this view are registered with Centralized Management.

Procedure

1. Log in with administrator credentials.

For example, `http://Frontend-hostname:58080/centralized-management`

2. Select **Licenses Management**.
3. In the **Licenses Listing** pane, read the **Expiration** column.

The *EMC Watch4net Centralized Management Guide* provides information about using Centralized Management and server registration.

Obtaining a new license after host reconfiguration

If you change the configuration of the host, for example, by adding interfaces, removing them, or by changing the operating system, your Watch4net permanent license may no longer work. In that case, obtain the host ID and request a new permanent license.

Before you begin

Make sure you have a login with root, APG, or system administrator privileges to perform this task. The user `apg` is the account that the application uses instead of root.

Procedure

1. Type the command for your operating system to obtain the host ID:

Operating system	Command
UNIX	<code>/opt/APG/bin/manage-licenses.sh host-id</code>
Windows	<code>C:\Program Files\APG\bin\manage-licenses.cmd host-id</code>

2. Copy the host ID information to a clipboard or to a text file.
3. For product licenses, go to `support.emc.com`.
4. Browse to **Service Center > Product Registration & Licenses**.
5. Paste the host ID information into the Service Request (SR) for a new permanent license.

CHAPTER 8

Ensuring Compliance

This chapter covers the following topics:

◆ Policy Management.....	94
◆ Notifications for compliance breaches.....	97
◆ EMC Support Matrix Management.....	99

Policy Management

You can use default policies or create user-defined policies to analyze configuration changes in your data center for compliance and to detect breaches against these policies.

- ◆ [Creating a policy on page 94](#)
- ◆ [Enabling or Disabling a policy on page 96](#)
- ◆ [Editing a policy on page 96](#)
- ◆ [Copying a policy on page 97](#)
- ◆ [Deleting a policy on page 97](#)

Creating a policy

Use the existing pre-defined templates to create a new policy or the empty policy template to customize it.

Procedure

1. Click **Operations**.
2. Under **Compliance**, click **Storage Compliance** > **Manage Rules & Policies**.
3. Click **Create Policy**.
4. Select a template from the **Select Policy from template** drop-down list.
5. Click **Description** and provide the following details:

Field	Description
Policy Name	Type a name for the policy.
Policy Description	Type a brief description for the policy.-
State	Select Enabled or Disabled from the drop-down list if you want the policy to run once it is created or otherwise.

Note

If you choose the state as **Disabled**, you should manually enable the policy after creating the policy to run it.

6. Click **Scope**.
 - a. To create a new user-defined scope, click **Create Scope**.
The **Create Scope** dialog box appears.
 - b. Provide the following details:

Field	Description
Name	Type a name for the scope.
Criteria	Type a criteria for the scope in the following format: <Property Name> <Operator> <Value>, such as; devtype = 'Array' & devicetype = 'FabricSwch' & parttype = 'port', and so on.

Note

Some examples of well-defined properties that can be used in criteria are:

- For Switch: devtype = 'Objecttype', device = 'SwitchName'.
 - For Fabric: parttype = 'Fabric/VSAN', part = 'FabricName'.
 - For Array: serialnb = 'Array S/N', ip = 'IP Address', revision = 'Microcode'.
 - For Host: domain = 'Domain Name', devdesc = 'Operating System Version', pmodel = 'HBA Model'.
-

Note

The **Criteria** allows the following operators: & (AND), I (OR), and= (Equal To). You can also use wild card entry, such as, % for starts with, ends with, and contains expressions.

- c. Click **Show Members** to validate the criteria.

The **Show Members** dialog box listing members with matching criteria appears. A warning message is displayed for invalid criteria.

- d. Click **Save**.

The new user-defined scope is listed in the Scope table.

Note

Click **Show Members** to view the list of members associated with the scope.

7. Click **Rules**.

- a. Click **Create Rule**.

- b. Select a rule name from the drop-down list.

- c. Double-click a rule to enable or disable it in the corresponding **State** column of the rule.

- d. Double-click a rule to set the **Severity** to **Critical**, **Major**, **Minor**, or **Normal** using the drop-down list.

- e. To provide criteria values for the rule, click **Criteria For Selected Rule** and select one of the following criteria and type a value in the corresponding text box:
 - Select the corresponding check box to enable the criteria or otherwise.
 - Select **LE** in the drop-down list to represent the Less than or Equal to value for the criteria.
 - Select **EQ** in the drop-down list to represent the Equal to value for the criteria.
 - Select **IN** in the drop-down list to represent the value(s) that the criteria contains.

Note

You can type multiple values that are comma separated. The values can contain different patterns and need not be restricted only to strings.

Note

Rules validate objects when its *vstatus* value is active. The other objects are ignored.

- f. Click **Recommendation** to view the recommendation for the rule.
8. Click **Schedule**.
 - a. Select a **Schedule** option and specify appropriate values from the drop-down lists to run the policy at a scheduled time.
9. Click **Save**.

Enabling or Disabling a policy

If you want the suite to run the policy after it is created, you must enable the policy. After enabling the policy, you can stop it any time by disabling it.

Procedure

1. Click **Operations**.
2. Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
3. Right-click on a policy and select either:
 - **Enable** - To enable the policy to run.
 - **Disable** - To disable the policy.

Editing a policy

You can modify a default or a user-defined policy to adhere to the configuration changes in your data center.

Procedure

1. Click **Operations**.
2. Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
3. Right-click on a policy and select **Edit**.
4. You can modify the following fields of a policy in the respective tabs:

Tab	Action(s)
Description	- In the Policy Name field, type a new name for the policy. - In the Policy Description field, type a new description for the policy. - Modify the State to enable or disable the policy to run it or otherwise.
Scope	Click  corresponding to a user-defined scope to delete it.
Rule	- Click Edit Rule to modify a rule. - Click Delete Rule to delete a rule.

Tab	Action(s)
	- Click Criteria for Selected Rule to modify a criteria value of a rule. - Click Recommendation to edit the recommendation description of a rule.
Schedule	Specify appropriate values using respective Schedule drop-down lists to run the policy at a different scheduled time.

- Click **Save**.

Copying a policy

If you want to retain the original default and user-defined policies, you can duplicate a policy into multiple copies and use them to analyze your data center for compliance.

Procedure

- Click **Operations**.
- Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
- Right-click on a policy and select **Copy**.

The **Copy Policy** dialog box stating **Are you sure you want to copy this policy?** appears.

- Click **OK**.

Deleting a policy

You can delete a default and a user-defined policy. However, a default policy can be added again as it is part of the default list of templates available in the system.

Procedure

- Click **Operations**.
- Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
- Right-click on a policy and select **Delete**.

The **Delete Policy** dialog box stating **Are you sure you want to delete this policy?** appears.

- Click **OK**.

Notifications for compliance breaches

The suite can generate email and SNMP trap notifications for compliance alerts. If you enable this feature, the suite forwards the compliance alerts generated as email messages or SNMP trap notifications to select members of your organization.

Default alert definitions allow you to set up this feature by using the default filters for the pre-defined policies and notification content as is. You have to set a new filter for new policies created to generate compliance alerts that will then be forwarded as email messages or SNMP traps notifications.

To generate compliance alerts and enable notifications you have to perform:

- [Configuring alert definitions for notification of compliance breaches on page 98](#)
- [Setting a filter for compliance breach notification on page 98](#)

3. [Enabling notifications for compliance breaches on page 99](#)

Configuring alert definitions for notification of compliance breaches

You must specify the recipient details for both email and trap notifications in the alert definitions before enabling them.

Before you begin

Ensure you have created a filter for any new compliance policies created to generate compliance alerts.

Procedure

1. Click **Administration**.
2. Under **Modules** area, click **Alerting**.
3. In the alerting front-end, click **Local Manager** > **Alert definitions** > **Compliance**.
4. Under **Alert definitions**, click **Compliance-Breach-Notification** > **Configure**.
5. Add the configuration details to the alert definition.

Alert definition	Mandatory configuration	Description
Mail Notification	To	The alert notification recipient's email address.
SNMP v1 Trap: Trap Notify template	Host	The trap destination IP address

Note

Email/SNMP notifications contain the following **State** values with their respective numbers: **Active** = 1 and **Resolved** = 0.

Note

Email/SNMP notifications contain the following **Severity** values with their respective numbers: **Critical** = 1, **Major** = 2, **Minor** = 3, and **Normal** = 5.

6. Click **Save**.

Setting a filter for compliance breach notification

The suite can generate compliance alerts for all the pre-defined policies in Storage Compliance. If you create a new policy for Storage Compliance, you have to create a filter for this new policy so that alerts generated due to this policy are collected.

Procedure

1. Click **Administration**.
2. Under **Modules** area, click **Alerting**.
3. In the alerting front-end, click **Local Manager** > **Alert definitions** > **Compliance**.
4. Click **Compliance-Breach-Notification**.
5. Double-click **Filtered entry Compliance Notification**.
6. In the **Edit Component** area, click **Filter**, select **or** condition, and type `policy name = " <policy name> "`.

7. Click **OK**.
8. Click **Save**.

Enabling notifications for compliance breaches

By default, all alert definitions for compliance alerts are in disabled state. After configuring the email destination or SNMP trap destination in the alert definition, you must enable it. The suite then starts forwarding all compliance alerts to the specified recipients.

Before you begin

Ensure you have created a filter for any new compliance policies created to generate compliance alerts.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the alerting front-end, click **Local Manager** › **Alert definitions** › **Compliance**.
4. Under **Alert definitions**, select the corresponding check box of an alert and click **Enable**.

Disabling notifications for compliance breaches

After enabling the generation of compliance alert notifications, you can disable it any time you choose. The suite stops forwarding alert notifications.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the alerting front-end, click **Local Manager** › **Alert definitions** › **Compliance**.
4. Under **Alert definitions**, select the corresponding check box of an alert and click **Disable**.

EMC Support Matrix Management

Use the EMC Support Matrix (ESM) data to compute interoperability information available in the data center. In addition, ESM Matching bridges the gap between discovered resources in the data center and the E-Lab data.

ESM matching is configured in the following ways:

- ◆ Auto-matching
- ◆ Multi-matching
- ◆ Manual-matching
- ◆ Bulk-matching

Note

Auto-matching and Multi-matching are configured as back-end processes. These processes match discovered resources data to a single or multiple configuration items, respectively, in the data center and E-Lab data, every time a new object is discovered.

Note

By default, if a configuration item has a discovered value which cannot be auto-matched, then the status of the configuration item appears as **Not Matched** in the corresponding **Matched Methods** column.

Manual-matching to ESM

If a configuration item has discovered values and is not automatically matched or multi-matched, then you can manually match the configuration item to the discovered values.

Procedure

1. Click **Operations**.
2. Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
3. Click  in the corresponding **Matched To** column of the configuration item with the status **Not Matched**.

The **Match Value** dialog box appears.

4. Type an appropriate value.
5. Click **OK**.

Results

The status of the configuration item appears as **Manual** in the corresponding **Matched Methods** column.

Bulk-matching to ESM

When multiple configuration items have the same discovered value, they cannot be auto-matched. In such a scenario, bulk-matching is enabled for all, when you manually match one of the configuration items.

Procedure

1. Click **Operations**.
2. Under **Compliance**, click **Storage Compliance** › **Manage Rules & Policies**.
3. Click  in the corresponding **Matched To** column of the configuration item with the status **Not Matched**.

The **Match Value** dialog box appears.

4. Type an appropriate value.
5. Click **OK**.

Results

The status of the configuration item appears as **Bulk** in the corresponding **Matched Methods** column.

CHAPTER 9

Alerting

This chapter contains the following topics:

- ◆ [Threshold alerts](#)..... 102
- ◆ [Availability alerts for PowerPath health metrics](#)..... 104
- ◆ [Alert notifications for consolidated alerts](#)..... 106
- ◆ [Alert retention](#)..... 112

Threshold alerts

Threshold alerts appear in the **Operations > Alerts > All Alerts** report and alerts **Dashboard**. The suite generates threshold alerts based on the metric data collected from discovered switches and storage systems. The threshold metrics set in the default alert definitions control the generation of threshold alerts.

Although the suite provides out-of-box support for the following SolutionPacks, you can use this feature to generate threshold metrics for any metric.

- ◆ SolutionPack for Brocade FC Switch
- ◆ SolutionPack for Cisco MDS/Nexus
- ◆ SolutionPack for EMC VNX
- ◆ SolutionPack for EMC VMAX

Managing threshold alerts involves:

- ◆ [Enabling threshold alerts on page 103](#)
- ◆ [Disabling threshold alerts on page 103](#)
- ◆ [Customizing a default threshold alert definition on page 103](#)
- ◆ [Creating a custom threshold alert definition for out-of-box SolutionPack support on page 104](#)

Prerequisites for generating threshold alerts

Use this topic as a checklist to verify if the requisite setup is complete for the suite to successfully generate threshold alerts based on data collected from external sources.

- ◆ Installation of the particular SolutionPack with the selection of the appropriate options for collection of metrics from the devices.
- ◆ Selection of these threshold metric consolidation options for out-of-box support:

Table 4 Out-of-box Alert Consolidation Options During SolutionPack Installation

SolutionPack	Option for consolidation of threshold metrics
SolutionPack for Brocade FC Switch	Enable the Brocade Port Link Utilization Alert
SolutionPack for Cisco MDS/Nexus	Enable the Cisco Port link Utilization Alert
SolutionPack for EMC VMAX	Enable the EMC VMAX Front-End Director Utilization Alert
	Enable the EMC VMAX Front-End Director Port Utilization Alert
SolutionPack for EMC VNX	Enable the VNX Block SP Utilization
	Enable the VNX Block Dirty Pages

- ◆ Successful discovery of the device.
The *EMC Storage Resource Management Suite Installation and Configuration Guide* describes the setup required for successful device discovery in detail.
- ◆ Enabling of the default **<Source>Alert Threshold Definition** or your custom threshold alert definition. For example, if you choose to use the suite's default alert definition

for VNX threshold alerts, then you should enable **Enable the VNX Block SP Utilization** and **Enable the VNX Block Dirty Pages**.

Enabling threshold alerts

By default, all the out-of-box alert definitions for consolidated alerts are in disabled state. If you want the suite to generate threshold alerts based on metric data consolidated from external sources, you must enable the threshold alert definitions for each source. You can use the default alert definitions available, or create a custom definition and then enable it.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** > **Alert definitions**.
4. Click the **Source**.
For example, **Brocade FC switch** for Brocade switches.
5. In the **Alert definitions** table, click the appropriate alert definition.
6. Click **Enable**.

Disabling threshold alerts

After enabling the generation of threshold alerts based on the metric data consolidated from external sources, you can stop it any time.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** > **Alert definitions**.
4. Click the **Source**.
For example, **Brocade FC switch** for Brocade switches.
5. In the **Alert definitions** table, click the appropriate alert definition.
6. Click **Disable**.

Customizing an out-of-box threshold alert definition

You can customize the filters that the suite uses to generate threshold alerts based on the metric data collected.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local Manager** > **Alert definitions** > **Source**.
4. Click the default threshold alert definition to modify.
5. Double-click the required filter components and make suitable changes.

Note

Notifications fail if you modify the **adapterName** filter and sequence of the trap content in the **Filtered entry** and **SNMP Trap** action components, respectively.

6. Click **Save**.
7. Click **Save and enable**.
8. Click **OK**.

Creating a custom threshold alert definition for out-of-box SolutionPack support

In order to ensure that you have all the necessary alert workflow in your custom threshold alert definition, copy the default action components from the out-of-box alert definitions.

Procedure

1. Click **Administration**.
 2. Under **Modules**, click **Alerting**.
 3. In **Alerting Frontend**, click **Local Manager** › **Alert definitions** › **«Source»**.
 4. Click the appropriate alert definition, and then click .
 5. Click the **«Source»** to which the alert definition should apply.
 6. Click  and type a suitable description for your alert definition.
 7. Click **OK**.
 8. Double-click the required filter components and make suitable changes.
-

Note

Notifications fail if you modify the **adapterName** filter and sequence of the trap content in the **Filtered entry** and **SNMP Trap** action components, respectively.

9. If you need to add new components, drag and drop them from the component list beside the whiteboard.
10. Add appropriate filters and link the component connectors to other relevant components.
11. Click **Save**.
12. Click **Save and enable**.
13. Click **OK**.

Results

The new alert definition appears in the definitions list for the source.

Availability alerts for PowerPath health metrics

You can manage the suite consolidation and display of EMC PowerPath® health alerts — autostandby path, path availability, and LUN availability — for SolutionPack for Physical Hosts and SolutionPack for VMware vCenter.

The suite can consolidate and display PowerPath autostandby path, path availability, and LUN availability alerts for IBM AIX, Sun Solaris, Windows, RHEL, HP-UX, and ESX hosts. These alerts appear in the **Operations** › **Alerts** › **All Alerts** report and alerts **Dashboard**.

Managing threshold alerts involves:

- ◆ [Disabling ESX host availability alerts for PowerPath health metrics on page 106](#)
- ◆ [Enabling ESX availability alerts for PowerPath health metrics on page 106](#)
- ◆ [Disabling physical host availability alerts for PowerPath health metrics on page 105](#)
- ◆ [Enabling physical host availability alerts for PowerPath health metrics on page 105](#)
- ◆ [Tasks for enabling alert notifications on page 107](#)
- ◆ [Disabling email alert notifications for consolidated alerts on page 109](#)
- ◆ [Disabling SNMP trap alert notifications for consolidated alerts on page 110](#)

Prerequisites for generating availability alerts

Based on PowerPath health metrics (autostandby path, path availability, and LUN availability), the suite generates availability alerts for IBM AIX, Sun Solaris, Windows, RHEL, HP-UX, and ESX hosts. Use this topic as a checklist to verify if the requisite setup is complete.

- ◆ Installation of the following tools on hosts:
 - Rtools and PowerPath/VE remote CLI (rpowermt) on the ESX hosts.
 - PowerPath CLI (powermt) on physical hosts.
The *EMC Storage Resource Management Suite Installation and Configuration Guide* provides details on the installation of each of these tools.
- ◆ Installation of PowerPath CLI (powermt) on physical hosts.
- ◆ Installation of SolutionPack for Physical Hosts and SolutionPack for VMware vCenter, with **Enable the Host PowerPath Alerts** selected.
- ◆ Enabling of the default **Physical Hosts** and **VMware vCenter** alert definitions.

Note

You cannot modify the default alert definitions, or create custom alert definitions to generate alerts for PowerPath health metrics.

Disabling physical host availability alerts

You can disable the alert definitions any time after SolutionPack for Physical Hosts installation. The suite immediately stops generating the alerts.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** > **Alert definitions** > **Physical Hosts**.
4. In **Alert definitions** table, select the appropriate alert definition.
5. Click **Disable**.

Enabling physical host availability alerts

You can enable the required alert definitions if they were disabled after SolutionPack for Physical Hosts installation. The alert definitions are enabled by default, if **Enable the Host PowerPath Alerts** is selected while installing SolutionPack for Physical Hosts.

Procedure

1. Click **Administration**.

2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** › **Alert definitions** › **Physical Hosts**.
4. In the **Alert definitions** table, select the appropriate alert definition.
5. Click **Enable**.

Disabling ESX host availability alerts

You can disable the alert definitions any time after SolutionPack for VMware vCenter installation. The suite immediately stops generating the alerts.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** › **Alert definitions** › **VMware vCenter**.
4. In the **Alert definitions** table, select the appropriate alert definition.
5. Click **Disable**.

Enabling ESX host availability alerts

You can enable the required alert definitions if they were disabled after SolutionPack for VMware vCenter installation. The alert definitions are enabled by default, if **Enable the Host PowerPath Alerts** is selected while installing SolutionPack for VMware vCenter.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend**, click **Local Manager** › **Alert definitions** › **VMware vCenter**.
4. In the **Alert definitions** table, select the appropriate alert definition.
5. Click **Enable**.

Alert notifications for consolidated alerts

The suite can generate email and SNMP v1 trap notifications for all alerts consolidated from external sources.

If you enable this feature, the suite forwards all availability alerts received from or performance-based threshold alerts generated for external sources as email messages or SNMP v1 trap notifications to select members of your organization and third party applications (such as a ticketing system).

Out-of-box global alert definitions allow you to quickly set up this feature by using the default filters and notification content as is, or by tailoring them to your requirements:

Alert definition	Default filter for generating notification
Alert Consolidation Email Notification	Alert severity is critical, major or minor.
Alert Consolidation Trap Notification	All alert severities, except normal.

Prerequisites for generating alert notifications

Use this topic as a checklist to verify if the suite is ready to successfully forward consolidated alerts as email and SNMP v1 trap notifications to specific members of your organization or third party applications, respectively.

- ◆ In **Administration** › **GLOBAL SETTINGS**, the SMTP server details are set.
- ◆ In **Centralized Management** › **SolutionPacks** › **Independent SolutionPacks**, the Alert-Consolidation (Primary Backend) is reconfigured.
- ◆ In the **Alerting Frontend** › **Local Manager** › **Alert definitions**:
 - The email and SNMP trap destinations are configured in the relevant alert definitions.
 - Email and SNMP alert notification definitions are enabled.
- ◆ In the **Alerting Frontend** › **Local Manager** › **Adapters**, **Alert Consolidation Adapter** is enabled.

Tasks for enabling alert notification

Perform these tasks to forward consolidated alerts as alert notifications.

- ◆ [Reconfiguring the Alert-Consolidation SolutionPackBlock on page 107](#)
- ◆ Enabled by default — [Enabling the alert consolidation adapter for consolidated alert notifications on page 107](#)
- ◆ [Configuring the alert definitions for notification of consolidated alerts on page 108](#)
- ◆ Enable the out-of-box alert definitions, as required:
 - [Enabling email alert notifications for consolidated alerts on page 109](#)
 - [Enabling SNMP trap alert notifications for consolidated alerts on page 109](#)

Reconfiguring the Alert-Consolidation SolutionPackBlock

You must reconfigure the Alert-Consolidation SolutionPackBlock for the suite to successfully forward alert notifications as emails and SNMP traps. This procedure updates the appropriate configuration files and restarts the Alerting service.

Procedure

1. Click **Administration** › **Centralized Management** › **SolutionPacks** › **Independent SolutionPackBlocks**.
2. In the Alert-Consolidation (Primary Backend) instance row, click .
3. In the **SolutionPack Reconfiguration** window that opens, select **Enable Alert Notification**.
4. Edit the **Host** and **Port** fields, if required.
5. Click **Reconfigure**.

Enabling the Alert Consolidation Adapter for consolidated alert notifications

When you enable the relevant alert definitions, the **Alert Consolidation Adapter** must be in the default enabled state to allow the suite to successfully forward consolidated alerts

as email or SNMP trap notifications. For some reason, if the adapter is disabled, enable it.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local ManagerAdapters**.
4. In the **Adapters** table, click **Alert Consolidation Adapter**.

Note

If you modify the pre-defined properties and port values of the adapter, notifications fail

5. Click **Enable**.

Configuring the alert definitions for notification of consolidated alerts

You must specify the recipient details for both email and trap notifications in the alert definitions before enabling them.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local Manager > Alert definitions**.
4. Click one these out-of-box alert definitions for notifications or your custom notification definition:
 - **Alert Consolidation Email Notification**
 - **Alert Consolidation Trap Notification**
5. Double-click the action components and add the configuration details:

Table 5 Mandatory configurations for email and trap alert notifications

Alert definition	Action component	Mandatory configuration	Description
Alert Consolidation Email Notification	Mail	To	The alert notification recipient's email address. The email address is set in Administration > GLOBAL SETTINGS .
Alert Consolidation Trap Notification	SNMP Trap	Host	The trap destination IP address.
		Port	The default port is 162. You can modify it, if required.

Note

Notifications fail if you modify the **adapterName** filter and sequence of the trap content in the **Filtered entry** and **SNMP Trap** action components, respectively.

6. Click **Test action** to ensure that you have specified the right information.
7. Click **Save**.

8. Click **Save and enable**.
9. Click **OK**.

Enabling email notifications for consolidated alerts

By default, alert definitions for consolidated alerts are in disabled state. After configuring the email destination in the alert definition, you must enable it. The suite immediately starts forwarding email notifications for consolidated alerts to the specified recipients.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In **Alerting Frontend**, click **Local Manager** › **Alert definitions**.
4. In the **Alert definitions** table, click **Alert Consolidation Email Notification**, or your custom alert definition.
5. Click **Enable**.

Enabling SNMP trap notifications for consolidated alerts

By default, alert definitions for consolidated alerts are in disabled state. After configuring the SNMP trap destination in the alert definition, you must enable it. The suite immediately starts forwarding SNMP v1 trap notifications for the consolidated alerts to the specified trap destination.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In **Alerting Frontend**, click **Local Manager** › **Alert definitions**.
4. In the **Alert definitions** table, click **Alert Consolidation Trap Notification**, or your custom alert definition.
5. Click **Enable**.

Disabling email notifications for consolidated alerts

After enabling the generation of email notifications for consolidated alerts, you can disable it any time you choose. The suite immediately stops forwarding consolidated alerts as email notifications.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local Manager** › **Alert definitions**.
4. In the **Alert definitions** table, click **Alert Consolidation Email Notification**, or your custom alert definition.
5. Click **Disable**.

Disabling SNMP trap notifications for consolidated alerts

After enabling the generation of SNMP trap notifications for consolidated alerts, you can disable it at any time. The suite immediately stops forwarding consolidated alerts as SNMP v1 trap notifications.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local ManagerAlert definitions**.
4. In the **Alert definitions** table, click **Alert Consolidation Trap Notification**, or your custom alert definition.
5. Click **Disable**.

Customizing an out-of-box alert definition for notification of consolidated alerts

You can customize the filters the out-of-box alert definitions, **Alert Consolidation Email Notification** and **Alert Consolidation Trap Notification**, that the SRM Suite uses to generate notifications for consolidated alerts.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local Manager > Alert definitions**.
4. Click one of these default alert definitions:
 - **Alert Consolidation Email Notification**
 - **Alert Consolidation Trap Notification**
5. Click **Edit**.
6. In the alert definition whiteboard, double-click the action component and make suitable changes to the action component, including these mandatory configuration details:

Table 6 Mandatory configurations for email and trap alert notifications

Alert definition	Action component	Mandatory configuration	Description
Alert Consolidation Email Notification	Mail	To	The alert notification recipient's email address. The email address is set in Administration > GLOBAL SETTINGS .
Alert Consolidation Trap Notification	SNMP Trap	Host	The trap destination IP address.
		Port	The default port is 162. You can modify it, if required.

Note

Notifications fail if you modify the **adapterName** filter and sequence of the trap content in the **Filtered entry** and **SNMP Trap** action components, respectively.

7. Click **Test action** to ensure that you have specified the right information.
8. Click **Save**.
9. Click **Save and enable** to immediately start forwarding notifications.
10. Click **OK**.

Creating a custom alert definition for notification of consolidated alerts

In order to ensure that you have all the necessary alert workflow in your custom alert definition for notifications, copy the default action components from the out-of-box alert definitions.

Procedure

1. Click **Administration**.
2. Under the **Modules** area, click **Alerting**.
3. In the **Alerting Frontend** window, click **Local Manager** > **Alert definitions**.
4. Click one of these out-of-box notification alert definition, and then click :
 - **Alert Consolidation Email Notification**
 - **Alert Consolidation Trap Notification**
5. Click the appropriate location in **Local Manager** > **Alert definitions**, and click .
6. Type a suitable description for your alert definition.
7. Click **OK**.
8. Double-click the required filter components and make suitable changes.

Table 7 Mandatory configurations for email and trap alert notifications

Alert definition	Action component	Mandatory configuration	Description
Alert Consolidation Email Notification	Mail	To	The alert notification recipient's email address. The email address is set in Administration > GLOBAL SETTINGS .
Alert Consolidation Trap Notification	SNMP Trap	Host	The trap destination IP address.
		Port	The default port is 162. You can modify it, if required.

Note

Notifications fail if you modify the **adapterName** filter and sequence of the trap content in the **Filtered entry** and **SNMP Trap** action components, respectively.

9. If you need to add new components, drag and drop them from the component list beside the whiteboard.

10. Add appropriate filters and link the component connectors to other relevant components.

11. Click **Save**.

12. Click **Save and enable**.

13. Click **OK**.

Results

The new alert definition appears in the definitions list.

Alert retention

All consolidated alerts are available only for a pre-defined period of 30 days from the date of creation.

At the end of the retention period, the suite deletes alerts older than 30 days.

CHAPTER 10

Online updates

This chapter contains the following topics:

- ◆ [Configuring server settings for online updates](#)..... 114
- ◆ [Enabling the online update task](#)..... 114
- ◆ [Disabling the online update task](#)..... 115
- ◆ [Executing online update task manually](#)..... 115

Configuring server settings for online updates

The Online Update module automatically downloads the latest versions of the SolutionPacks from the EMC Update server to the module repository. It eliminates the need for manual check and downloads of the latest versions of SolutionPacks and modules. To receive online updates to the software, do the following:

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. Click **Online Update**.
4. Provide the following server setting details in the **Online Update Server Settings** page:
 - Update Server URL

Note

To connect to EMC Update server, type `https://vupdate.emc.com/SRM/PackageUpdates`.

- Username
 - Password
5. Click **Save**.

Enabling the online update task

Enable Online Update task to download the latest versions of the SolutionPacks and modules from the EMC Update server automatically.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. In the **Physical Overview** page, click **<host_name> - Front End** on which the Online Update task needs to be executed.
4. Click **Tasks**.
5. Type **OnlineUpdate** in the Search bar.
6. Click the **OnlineUpdate** scheduled task.
7. Click **Enable**.

Note

By default, this task is set to run once everyday at 12AM. You can customize the task schedule by editing the configuration file. The *EMC APG - Task Scheduler User Guide* provides more information on editing the task schedule.

Disabling the online update task

Disable Online Update task to manually download the latest versions of the SolutionPacks and modules from the EMC Update server.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. In the **Physical Overview** page, click *<host_name>* - **Front End** on which the Online Update task needs to be executed.
4. Click **Tasks**.
5. Type **OnlineUpdate** in the Search bar.
6. Click the **OnlineUpdate** scheduled task.
7. Click **Disable**.

Executing online update task manually

At any time, you can run the online update task manually to access available software updates.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. In the **Physical Overview** page, click *<host_name>* - **Front End** on which the Online Update task needs to be executed.
4. Click **Tasks**.
5. Type **OnlineUpdate** in the Search bar.
6. Click the **OnlineUpdate** scheduled task.
7. Click **Run Now**.

CHAPTER 11

Troubleshooting

This chapter includes the following topics:

◆ SSL certificates	118
◆ Switch discovery failure in the SolutionPack for Cisco MDS/Nexus	121
◆ Troubleshooting service start-up problems on Windows	122
◆ Troubleshooting service start-up problems on UNIX	122
◆ Troubleshooting collector errors with Collector-Manager log files	123
◆ Licenses have expired	123
◆ Collect diagnostics on UNIX	124
◆ Collect diagnostics on Windows	125
◆ Troubleshooting alert collection errors on SolutionPack for EMC Data Protection Advisor	125
◆ Troubleshooting report generation errors on SolutionPack for EMC Data Protection Advisor	126

SSL certificates

Generally, EMC products are distributed with a self-signed SSL certificate. If you receive an SSL certificate error or a blank pane when you are attempting to access content served by a data source, resolve the issue by doing one of the following procedures.

[SSL certificate error examples on page 118](#) provides a list of sample errors.

Temporary resolution

Accept the certificate exception temporarily: [Troubleshooting an SSL certificate error: Accepting the exception on page 119](#).

Permanent resolution

Add an SSL certificate to your truststore to permanently prevent the exceptions from appearing.

You can:	Perform:
Install the self-signed SSL certificate that is provided with the EMC product using an Internet Explorer browser.	Installing a self-signed SSL certificate from an Internet Explorer browser on page 119
Install the self-signed SSL certificate that is provided with the EMC product using a Chrome or Firefox browser.	1. Export the self-signed SSL certificate to save it to the export location: Exporting an SSL certificate from an Internet Explorer browser on page 120 2. Importing an SSL certificate into your truststore on page 119
Import an exported SSL certificate that you or another administrator prepared.	Importing an SSL certificate into your truststore on page 119

If you need to prepare an SSL certificate to share with other users, use: [Exporting an SSL certificate from an Internet Explorer browser on page 120](#).

SSL certificate error examples

The SSL certificate error or a blank pane appears because the certificate is self-signed and it is not added as an exception or into the truststore yet. For example, if you are attempting to access content served by a data source, a blank pane appears when you select views under **SolutionPacks**.

Chrome error examples

This site's security certificate is not trusted!

Or, a blank pane.

Firefox error example

This Connection is Untrusted

Internet Explorer error examples

There is a problem with this website's security certificate.

Content was blocked because it was not signed by a valid security certificate.

Troubleshooting an SSL certificate error: Accepting the exception

An SSL certificate error appears in your browser.

To resolve this issue,

Procedure

- ◆ Accept the exception.

Steps may vary for your browser. For example, for Internet Explorer, click **Continue to this website (not recommended)**.

This temporary solution causes the exception to be ignored for the current browser session, but the exception reappears if you open another browser tab or window.

Installing a self-signed SSL certificate from an Internet Explorer browser

Install the self-signed SSL certificate that is provided with the EMC product to prevent certificate exception errors.

The following procedure is for Internet Explorer version 9.0. The steps may vary slightly depending on your browser version.

Procedure

1. Click **Continue to this website (not recommended)** in the browser window.
2. Click the red **Certificate Error** indicator next to the address bar.
3. In the **Untrusted Certificate** window, select **View certificates**.
4. In the **Certificate** dialog box, in the **General** tab, click **Install Certificate**.
5. Follow the instructions in the **Certificate Import Wizard**.
 - a. Select the **Place all certificates in the following store** setting to add the certificate to your truststore.
6. In the **Certificate** dialog box, in the **General** tab, click **OK**.

After you finish

Refresh the browser window.

Importing an SSL certificate into your truststore

Import an SSL certificate into your truststore to prevent certificate exception errors.

Procedure

- ◆ Import the certificate using one of the following procedures:
 - [Importing an SSL certificate in Chrome on page 119](#)
 - [Importing an SSL certificate in Firefox on page 120](#)
 - [Importing an SSL certificate in Internet Explorer on page 120](#)

Importing an SSL certificate in Chrome

For Chrome, use this procedure to import the certificate into your truststore.

The following procedure is for Chrome version 28.0.1500.95. The steps may vary slightly depending on your browser version.

Procedure

1. Go to **Settings > Show advanced settings**.
2. Scroll down and select **HTTP/SSL > Manage Certificates**.
3. In the **Trusted Root Certification Authorities** tab, click **Import**.
4. Follow the instructions in the **Certificate Import Wizard** to import the certificate.
5. Click **OK** in **Internet options**.
6. Refresh the browser window.

Importing an SSL certificate in Firefox

For Firefox, use this procedure to import the certificate into your truststore.

The following procedure is for Firefox version 23.0. The steps may vary slightly depending on your browser version.

Procedure

1. Go to **Firefox > Options > Options**.
2. Click **Advanced**.
3. In the **Certificates** tab, click **View Certificates**.
4. Click **Import**.
5. Locate the certificate file and click **Open**.
6. Click **Save**.
7. Click **OK** in **Certificate Manager**.
8. Refresh the browser window.

Importing an SSL certificate in Internet Explorer

For Internet Explorer, use this procedure to import the certificate into your truststore.

The following procedure is for Internet Explorer version 9.0. The steps may vary slightly depending on your browser version.

Procedure

1. Go to **Tools > Internet options**.
2. In the **Content** tab, click **Certificates**.
3. Select the **Trusted Root Certification Authorities** tab.
4. Click **Import**.
5. Follow the instructions in the **Certificate Import Wizard** to import the certificate.
6. Click **OK** in **Internet options**.
7. Refresh the browser window.

Exporting an SSL certificate from an Internet Explorer browser

Use this procedure to export a self-signed SSL certificate, so that you can either import the resulting exported certificate into your truststore or share it with other users.

The following procedure is for Internet Explorer version 10.0. The steps may vary slightly depending on your browser version.

Procedure

1. Access the host for which you want to export the SSL certificate from the browser.
For example: `https://<myHostname>.emc.com`
2. Select **Continue to this website (not recommended)** in the **There is a problem with this website's security certificate** window.
3. Click the red **Certificate Error** indicator next to the address bar.
4. In the **Untrusted Certificate** window, select **View certificates**.
5. In the **Certificate** dialog box, select the **Details** tab and ensure that the Show field is set to All. Click **Copy to File**.
6. In the **Certificate Export Wizard**,
 - a. Click **Next** in the **Welcome** screen.
 - b. Click **Next** to accept the default DER encoded binary X.509 setting in the **Export File Format** screen.
 - c. Browse to the location where you want to save the certificate.
 - d. Specify the certificate file name and click **Save**.
 - e. In the **File name** field, ensure that the location and the certificate file name are correct. Click **Next**.
 - f. Click **Finish**.
 - g. Click **OK** when The export was successful appears.
7. In the **Certificate** dialog box, click **OK** in the **Details** tab.
8. Share the resulting exported SSL certificate with other users to use.


Switch discovery failure in the SolutionPack for Cisco MDS/Nexus

Edit and exclude the **GENERIC-INTERFACES** capability on Cisco MDS/Nexus switches.

Discovery must be successful for every switch that the SolutionPack for Cisco MDS/Nexus displays. The value for the **devtype** must be **FabricSwitch**. For most switches, this information is discovered automatically. However, if your switch discovery does not occur with **devtype** equal to **FabricSwitch**, the switch will not appear under the SolutionPack reports.

Procedure

1. Edit device settings.
Open the Device Discovery GUI (`http://yourserver_ip:58080/device-discovery/`).
2. Exclude the **GENERIC-INTERFACES** capability on Cisco MDS/Nexus switches.
Without it, more polling occurs.



3. To exclude the capability, edit SNMP agent settings (under **Devices**), and exclude it.

Troubleshooting service start-up problems on Windows

Check the log files when services do not start.

Before you begin

Make sure you have logged in with system administrator credentials to check the log files.

Procedure

- ◆ Check these logs in the C:\Program Files\APG directory to troubleshoot start-up problems:

```
Databases\MySQL\Default\data\[SERVER NAME].err.  
Backends\Alerting-Backend\Default\logs>alerting-0-0.log  
Backends\APG-Backend\Default\logs\cache-0-0.log  
Collecting\Collector-Manager\Default\logs\collecting-0-0.log  
Web-Servers\Tomcat\Default\logs\service.log  
Tools\Task-Scheduler\Default\logs\scheduler-0-0.log  
Tools\Webservice-Gateway\Default\logs\gateway-0-0.log
```

Troubleshooting service start-up problems on UNIX

Check the log files when services do not start.

Before you begin

Make sure you have logged in with root or APG privileges to check the log files.

Procedure

- ◆ Check these logs in the /opt/APG/ directory to troubleshoot start-up problems:

```
Databases/MySQL/Default/data/[SERVER NAME].err  
Backends/Alerting-Backend/Default/logs/alerting-0-0.log  
Backends/APG-Backend/Default/logs/cache-0-0.log  
Collecting/Collector-Manager/Default/logs/collecting-0-0.log  
Web-Servers/Tomcat/Default/logs/service.log  
Tools/Task-Scheduler/Default/logs/scheduler-0-0.log  
Tools/Webservice-Gateway/Default/logs/gateway-0-0.log
```

Troubleshooting collector errors with Collector-Manager log files

Review the Collector-Manager log files to troubleshoot problems with data collection.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. Expand **Collecting**.
4. Click the Collector-Manager for your collector instance.

Collector-Manager::*<Collector-Manager instance>* - *<physical_host_ID>*

5. Expand **Log Files** and click the **View File** icon to review the configuration error messages.

Licenses have expired

If all licenses expire, you are unable to access the UI. Download and install a license file on the frontend server, upload the licenses to the product, and restart the tomcat service. It is only necessary to restart the tomcat service when the license for a web application, such as Web Portal, expires.

Installing licenses

Before the 30-day trial licenses expire, contact EMC Online Support to obtain and install new license files.

Before you begin

- ◆ Obtain the required license file by creating a Service Request on EMC Online Support.
- ◆ Make sure you have a login with root, APG, or system administrator privileges to perform this task. The user apg is the account that the Watch4net application uses instead of root.

Procedure

- ◆ Type the command for your operating system from the `bin` directory:

Operating system	Command
UNIX	<code>./manage-licenses.sh install <download_path>/<license.zip></code>
Windows	<code>manage-licenses.cmd install <download_path>\<license.zip></code>

Linux example:

```
./manage-licenses.sh install /opt/APG/licenses/license.zip
```

A `.zip` file can contain any number of license files.

Adding licenses

You must upload a license file whenever you install a new SolutionPack or ReportPack. The license enables views and reports to display.

Before you begin

- ◆ Obtain the required license file by creating a Service Request on EMC Online Support.

Procedure

1. Log in to the host with administrator credentials.
For example, `http://Frontend-hostname:58080/APG`
2. Click **Administration**.
3. In the **Centralized Management** pane, select **Licenses**.
4. In the Physical Overview tree, select **Licenses Management**.
5. In the **Licenses Listing** pane, click **Upload**.
6. In the **License Upload** dialog box, browse to your license file. Click **OK**.
Licenses in a .zip file are extracted and appear in the **Licenses Listing** pane.
7. When **License Upload Complete** appears, click **OK**, then click **Continue**.
8. If you have remote servers in a distributed environment, click **Synchronize** to add the licenses to the remote servers.

Restarting the tomcat service

If the tomcat service stops, restart it.

Before starting core software services, ensure the library search path is correct on the server hosting core software. This requirement applies to the libstdc package on Solaris and the libaio package on Linux.

- ◆ Type the command for your operating system from the bin directory of the core software installation:

Operating system	Command
UNIX	<code>./manage-modules.sh service restart tomcat</code>
Windows	<code>manage-modules.cmd service restart tomcat</code>

Results

* Starting 'tomcat Default'... [OK]

Collect diagnostics on UNIX

You can run the `diagnostic.sh` script to collect data for Customer Support. The script collects files required for troubleshooting Watch4net and compresses them into a zip file.

Procedure

1. Log in to a UNIX machine with root access.
2. Access the `diagnostic.sh` script at `/opt/APG/Tools/APG-Diagnostic-Tools/Default/bin/diagnostic.sh` and run it.

If Watch4net is installed as Minimal, the script is missing from the installation and you need to run `# manage-modules.sh install diagnostic-tools` to install the script.

Output from the script goes into a zip file in the `/tmp/` directory. The output file name is `APG-DiagnosticFiles.tar.gz`

3. Upload the output file into a Service Request or if the file is large, to an FTP site specified by Customer Support.

Collect diagnostics on Windows

You can run the `diagnostic.sh` script to collect data for Customer Support. The script collects files required for troubleshooting Watch4net and compresses them into a zip file.

Procedure

1. Log into a Windows machine and open a command window.
2. Navigate to the folder `<%INSTALL_ROOT%>\Program Files\APG\bin.`

If Watch4net is installed as Minimal, the `diagnostic.sh` script is missing and you need to run `manage-modules.cmd install diagnostic-tools` to install the script.

3. Run `diagnostic.cmd`.

Output goes into the file `C:/Program Files/APG/APG-DiagnosticFiles.zip`.

4. Upload the file into a Service Request or if the file is large, to an FTP site specified by Customer Support.

Troubleshooting alert collection errors on SolutionPack for EMC Data Protection Advisor

Reconfigure the data collection options in the SolutionPack for EMC Data Protection Advisor if alert collection is not occurring. This procedure applies to each DPA instance for which alert collection is not occurring.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. Click **Storage > EMC Data Protection Advisor**.
4. Within the **Instance** column, click on the edit icon to reconfigure the options for the affected DPA instance that is affected by the problem.
5. Select the **Enable alert collection from DPA**.
6. Click **Reconfigure**.

Troubleshooting report generation errors on SolutionPack for EMC Data Protection Advisor

Reconfigure the data collection options in the SolutionPack for EMC Data Protection Advisor if reports are not being generated. This procedure applies to each DPA instance for which reports are not being generated.

Procedure

1. Click **Administration**.
2. Click **Centralized Management**.
3. Click **Storage > EMC Data Protection Advisor**.
4. Within the **Instance** column, click on the edit icon to reconfigure options for the affected DPA instance.
5. Select the **Do you want to configure advanced settings** and **Enable scheduling of DPA reports** options.
6. Click **Reconfigure**.

CHAPTER 12

Reference

This chapter includes the following topics:

◆ Report URL syntax	128
◆ Uploading your company logo	131
◆ Importing users	132
◆ Configuring the user process limits for the Linux installation	132
◆ Logging in to the User Interface	132
◆ Starting the core software services from the command line	133
◆ Starting the core software services on Windows	133
◆ Verifying if the core software services are running	134
◆ Connecting to Centralized Management	135
◆ Stopping core software services on a UNIX server	135
◆ Stopping core software services on a Windows server	135

Report URL syntax

The reports in the interface are available to third-party client tools that use the URL syntax and the auto-login functionality.

URLs

The base URL is: `http://[APGserverIP:port#]/APG/?param=value¶m=value`

The URL to display reports on a full page is: `http://[APGserverIP:port#]/APG/report.jsp?param=value¶m=value`

The URL to display reports in jpg, png, pdf, svg, xls, and csv formats is: `http://[APGserverIP:port#]/APG/report.format?param=value¶m=value`

You must replace *format* in the above URL with the desired format. For example: `http://[APGserverIP:port#]/APG/report.csv?param=value¶m=value`

Tree management

Syntax	Description
<code>select=<nodeid></code>	Displays the report of the specified node.
<code>expand=<nodeid></code>	Expands the selected node and then the parent nodes.
<code>collapase=<nodeid></code>	Collapses the selected node and any child nodes.
<code>collapaseALL</code>	Collapses the whole tree.

Report settings

Syntax	Description
<code>display=<code></code>	The code numbers for the display modes: ◆ 0=normal display mode ◆ 1=summary display mode with one graph per line ◆ 2=summary display mode with two graphs per line ◆ 3=summary display mode with three graphs per line
<code>node=<code></code>	The codes for the report modes: ◆ srt=table report mode ◆ lrt=leaf table report mode ◆ frg=graph report mode ◆ srg=children report mode ◆ nrx=node report mode ◆ vrx=baseline report mode ◆ stk=stacked chart report mode ◆ hb=horizontal bar chart report mode ◆ pie=pie chart report mode

	◆ gg=gauge chart report mode ◆ ico=icon report mode ◆ mix=basic mixed report mode ◆ dmx=mixed default preferences report mode
period=<seconds>	The period, in seconds, of the selected aggregate. Zero indicates real time. This period must exist in the database.
type=<code>	Not applicable when the period is zero. The code numbers for the aggregate types are: ◆ 3=average ◆ 4=minimum ◆ 5=maximum ◆ 6=sum ◆ 7=last value ◆ 8=number of aggregated values ◆ 9=timestamp of the last aggregated value
var_idx=<id>_<id>_<id>...	A list of indexes separated by underscores that restricts the displayed variables.

Graph display preferences

Syntax	Description
width=<pixels>	The graph width in pixels.
lower=<value>	The lowest value displayed on graphs.
upper=<value>	The highest value displayed on graphs.

Time management

Syntax	Description
durationType=<code>	The code to use as the default for the report time range description: ◆ n=not applicable because duration is based on something else such as a timestamp ◆ p=previous duration ◆ l=last duration ◆ c=current duration
duration=<code>	The code to use as the default for the report time range:

	◆ a=custom code that is a combination of numbers and units ◆ s=seconds ◆ m=minutes ◆ h=hour ◆ d=day ◆ w=week ◆ M=month ◆ y=year You can use as many duration codes as needed for the report time range but you cannot duplicate a code. You can also use calendar in the duration code followed by start and end or start_ts and end_ts, as explained in the following rows.
start=<date>end=<date>	The time range of the requested report. The <date> field has the YYYY-MM-DD HH:MM format. Note URL encoding may transform spaces into plus signs and colons and periods into %3A.
start_ts=<timestamp>end_ts=<timestamp>	The time range using timestamps of the requested report. These are UNIX timestamps.

Advanced time management

Syntax	Description
itz=<timezone code>	The time zone for the report, which uses standard zone names such as America/Montreal.
tf=<time filter.expression>	The maintenance period for the report. The <time filter.expression> field consists of <dom>;<dow>;<hod> to represent the following: ◆ dom=comma-separated list of days in a month between 1 and 31. For example, 1, 2, 10 represent the first, second, and tenth day of the month. ◆ dow=comma-separated list of days in a week between 1 and 7, with 1 starting on Sunday. For example, 1 and 2 represent Sunday and Monday.

	♦ <code>hod</code>=comma-separated list of hours in a day between 0 and 23. For example, 12 and 13 represent 12pm and 1pm.
--	--

Query and variable selection

When using this syntax, always check that the URL encoding did not change spaces into plus signs.

Syntax	Description
<code>search-base=<base></code>	The node from which the search begins. The <i><base></i> field can contain the root node, which is the default, or another node combined with a parameter.
<code>q=<search string></code>	The query string for the standard or quick search mode. It is a space-separated list of tokens.
<code>qsp=<search properties></code>	The space-separated list of properties in which to search.
<code>qsg=<grouping properties></code>	The space-separated list of node expansion to group query results.
<code>qf=<search filter></code>	The variable selection filter.
<code>qg=<grouping properties></code>	The space-separated list of node expansion to group query results.

Auto-login

Syntax	Description
<code>autologin=<user>:<pass></code>	Bypasses the Login page by supplying user credentials to the URL. Since the colon is part of the auto-login syntax, the username and password cannot contain a colon.

Uploading your company logo

You can replace the placeholder in the **Login** page with your company logo. The default logo dimension is 232 pixels wide by 63 pixels high.

Procedure

1. Log in with administrator credentials for core software.
2. Click **Administration**.
3. Under **Frontend Management**, click **Portal > Default Display**.
4. Click **Add a New Logo**, browse to the logo, and click **OK**.
5. Select your newly uploaded logo in the **Logo** list box and click **Save**.

Importing users

You can set up users individually in the User Interface or import users with the Administration Tool.

Procedure

- ◆ For instructions, access the `APG-Administration-Tool.pdf` in the documentation directory of the installation path. For example, `/opt/APG/Doc`.

Configuring the user process limits for the Linux installation

Increase the user process limits for the user account to a maximum of 65534. This modification enables services to open 65534 files and 65534 processes when needed.

Before you begin

- ◆ Make sure you have a login with root privileges.
- ◆ The core software installed on a server running Red Hat Enterprise Linux 6, CentOS Linux 6, or SUSE Linux Enterprise Server (SLES) 11.
- ◆ The name specified for the user account during the installation.

Procedure

1. Open the `/etc/security/limits.conf` file.
2. Add the following lines for the user.

In this example, the user is `apg`.

```
apg    hard    nofile   65534
apg    soft    nofile   65534
apg    hard    nproc    65534
apg    soft    nproc    65534
```

3. Save the file.
4. Type the following commands from the `bin` directory of the installation to restart the services:

```
./manage-modules.sh service stop all
./manage-modules.sh service start all
```

5. Type the following command to verify the changes:

```
su apg -c 'ulimit -n -u'
```

Logging in to the User Interface

Log in to continue to view reports or to test user access by specifying operator credentials.

Procedure

1. Point your browser to the URL.

For example, `http://<Frontend-hostname>:58080/APG`

2. Log in.

- a. Type the login credentials.

The default username is **admin** and the default password is **changeme**.

- b. Click **Sign In**.

After you finish

Note

You are automatically logged off after four hours.

Starting the core software services from the command line

After installing core software, you must start all required services.

Before you begin

- ◆ Before starting core software services, ensure the library search path is correct on the server hosting core software. This requirement applies to the libaio package on Linux.
- ◆ Make sure you have a login with root, APG, or system administrator privileges to perform this task. The user apg is the account that the core software application uses instead of root.

Procedure

- ◆ Type the command for your operating system from the `bin` directory of the installation:

Operating system	Command
UNIX	<code>./manage-modules.sh service start all</code>
Windows	<code>manage-modules.cmd service start all</code>

Results

A list of services appears. For example:

```
* Starting 'topology-mapping-service Default'... [ OK ]
* Starting 'webservice-gateway Default'... [ OK ]
* Starting 'mysql Default'... [ OK ]
* Starting 'alerting-backend Default'... [ OK ]
* Starting 'backend Default'... [ OK ]
* Starting 'collector-manager Default'... [ OK ]
* Starting 'collector-manager Generic-SNMP'... [ OK ]
* Starting 'collector-manager Load-Balancer'... [ OK ]
* Starting 'collector-manager emc smarts'... [ OK ]
* Starting 'event-processing-manager emc smarts'... [ OK ]
* Starting 'tomcat Default'... [ OK ]
* Starting 'task-scheduler Default'... [ OK ]
```

Starting the core software services on Windows

After installing core software on Windows, you can start all required services from a Windows desktop.

Before you begin

Make sure you have logged in with system administrator credentials to manage services.

Procedure

1. From the application server's desktop, select **Start > Administrative Tools > Services**.
2. Right-click each service and click **Start**, in this order:

- APG Topology Mapping Service
- APG Topology Service
- APG Webservice Gateway
- APG MySQL
- APG Alerting Backend
- APG Backend
- APG Collector Manager (Generic RSC)
- APG Collector Manager (Generic-SNMP)
- APG Collector Manager (Load-Balancer)
- APG Event Processing Manager (Alert-Consolidation)
- APG Tomcat
- APG Task Scheduler

Verifying if the core software services are running

Verify the core software services are running by obtaining the status.

Before you begin

Make sure you have a login with root, APG, or system administrator privileges to perform this task. The user `apg` is the account that the core software application uses instead of root.

Procedure

1. Type the command for your operating system from the `bin` directory of the core software installation:

Operating system	Command
UNIX	<code>./manage-modules.sh service status all</code>
Windows	<code>manage-modules.cmd service status all</code>

2. Verify that each service has a status of running in the output:

Results

```
* Checking 'topology-mapping-service Default'... [ running ]
* Checking 'webservice-gateway Default'... [ running ]
* Checking 'mysql Default'... [ running ]
* Checking 'alerting-backend Default'... [ running ]
* Checking 'backend Default'... [ running ]
* Checking 'collector-manager Default'... [ running ]
* Checking 'collector-manager Generic-SNMP'... [ running ]
* Checking 'collector-manager Load-Balancer'... [ running ]
* Checking 'collector-manager emc smarts'... [ running ]
* Checking 'event-processing-manager emc smarts'... [ running ]
* Checking 'tomcat Default'... [ running ]
* Checking 'task-scheduler Default'... [ running ]
```

Connecting to Centralized Management

Connect to the server so that you can access Centralized Management to install and administer SolutionPacks.

Centralized Management is one of the multiple web applications available in core software.

Procedure

1. Open a browser.
2. Type `http://<serverIP:port>/centralized-management`.

The *serverIP* is the fully qualified hostname or IP address of the server where core software is running.

The *port* is the port number of the server.

For example:

```
http://myHost.emc.com:58080/centralized-management
```

3. Log in.
 - a. Default username is **admin**.
 - b. Default password is **changeme**.
 - c. Click **Sign In**.

After you finish

You are automatically logged off after four hours.

Stopping core software services on a UNIX server

Use the `manage-modules.sh service stop` command to stop a specific core software service or to stop all core software services on a UNIX server.

Before you begin

Make sure you have logged in with root or APG privileges.

Procedure

- ◆ Type `manage-modules.sh service stop <service_name>` from the `bin` directory of the installation to stop a specific core software service.

This example shows how to stop all core software services:

```
./manage-modules.sh service stop all
```

Stopping core software services on a Windows server

Use this procedure to stop core software services from the Windows desktop.

Before you begin

Make sure you have logged in with system administrator credentials to manage services.

Procedure

1. From the application desktop, select **Start > Administrative Tools > Services**.

2. Right-click each service and click **Stop**, in order:

- APG MySQL
- APG Backend
- APG Alerting Backend
- APG Collector Manager
- APG Tomcat
- APG Task Scheduler
- APG Webservice Gateway